

Technical Report 93-1

A Cognitive Framework for Battlefield Commanders' Situation Assessment

**Marvin S. Cohen, Leonard Adelman, Martin A. Tolcott,
Terry A. Bresnick, F. Freeman Marvin**

Cognitive Technologies, Inc.
4200 Lorcom Lane
Arlington, VA 22207

7 October 1993

**United States Army Research Institute
Fort Leavenworth Field Unit**

Technical Report 93-1

A Cognitive Framework for Battlefield Commanders' Situation Assessment

**Marvin S. Cohen, Leonard Adelman, Martin A. Tolcott,
Terry A. Bresnick, F. Freeman Marvin**

Cognitive Technologies, Inc.
4200 Lorcom Lane
Arlington, VA 22207

**United States Army Research Institute
Fort Leavenworth Field Unit
Fort Leavenworth, Kansas 66027-0347**

7 October 1993

The views, opinions, and findings contained in this report are those of the authors and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other official documentation

ACKNOWLEDGMENTS

We are grateful to Dr. Jon Fallesen of ARI for his advice and support throughout the preparation of this report. We also appreciate the comments and feedback of many others at the Fort Leavenworth ARI Field Unit. We wish to thank Captain Clay Miller for his work arranging our visits to four different Army posts. And we are especially grateful to the many officers who agreed to serve as subjects. Finally, thanks to General Leonard P. Wishart, III, USA (ret.) and General Charles P. Otstott, USA (ret.) for their invaluable advice at various stages of the project.

EXECUTIVE SUMMARY

Requirement:

Situation assessment is the basis for many of the planning activities performed by the battlefield commander and staff. Improved situation assessment may lead to faster, better planning. Yet the cognitive processes and skills involved in situation assessment are not as yet well understood.

Procedure:

A cognitive framework for battlefield commanders' situation assessment was developed. The framework draws on published work in cognitive psychology and in the tactical battlefield domain, as well as on interviews conducted with active-duty command staff. Components of the framework were selected for inclusion based on empirical support in the literature and relevance to situation assessment performance. The components were integrated into a comprehensive framework that specifies their relationships and interactions. Finally, aspects of the framework were illustrated with actual experiences of tactical battlefield staff. This framework is still in development, and will be refined and perhaps modified as more observational and experimental data are obtained.

Findings:

At the most general level, the framework has three principal components: memory structures, value/action structures, and metacognitive (or executive) processes. Memory structures range from the highly temporary contents of working memory (the attended part of the current situation model) to relatively permanent information in long-term memory. Between these two extremes is implicit focus (parts of the situation model which are readily available for attention though not currently attended) and current episodic memory (a history of efforts in the current problem).

Structural constraints on situation assessment performance are defined in terms of these components: i.e., limited space in working memory, time and errors in retrieval from long-term memory, and cognitive effort required by executive processes. Different modes of processing are likewise defined in terms of the framework components: Procedural processing involves a direct link between a situation model and activation of a response within working memory; knowledge-based behavior requires retrieval of knowledge from long-term memory and in some cases the construction of novel situation models. Intuitive behavior utilizes domain-specific knowledge structures developed over experience, while analytic behavior utilizes general-purpose knowledge structures learned from instruction.

A variety of long-term memory knowledge structures are important in situation assessment. Enemy *plan structures* organize information about enemy

goals, strengths, and opportunities, and describe how they lead to intentions, actions, and consequences. *Enemy goal structures* organize information about the trajectory of the enemy's major goals in time, high-level principles for achieving those goals, and specific actions which realize the goals in accordance with the principles. *Temporal plan execution* structures provide a more detailed description of the temporal durations, precedence relations, and contingencies among actions. *Enemy planning/C²* structures describe the enemy roles and activities involved in producing, communicating, and implementing plans. *Terrain* structures relate terrain features to expected enemy actions and prescribed friendly actions. Different situation assessors may *frame* problems differently, i.e., bring different knowledge to bear on it. Alternative frames include *proactive*, *predictive*, or *reactive* high-level principles, and conceptualizing enemy plans primarily in terms of goals, terrain, or strength.

Value/actions structures are not a separate compartment of knowledge; rather, they reflect a qualitatively different way of viewing knowledge. They represent the extent to which possible states of affairs are valued, as distinct from how strongly their existence is believed. Value/action structures, like memory structures, vary from the highly temporary (i.e., the currently executed portion of an action or plan), to the relatively permanent (i.e., high-level values and long-term goals). Between these two extremes fall the current plan and the trajectory image (a sequence of desired situations or major goals for the current problem or situation).

Metacognitive processes shape and guide the retrieval of knowledge from long-term memory and its synthesis in a model and/or plan for the current situation. Metacognition includes components of quick verification, full verification, and facilitation. Quick verification assesses the time available before a decision must be made, the stakes of the decision, and the degree of confidence in the current solution, and determines whether additional processing is required. Full verification tests the current situation model to determine whether it is incomplete, unreliable, or in conflict. Facilitation guides the collection of new data, revised interpretation of the current data, or the activation of additional knowledge from long-term memory to correct any problems that are found. Steps taken to resolve one kind of problem can produce other problems. For example, conflicting data can be fit into a single coherent situation model (or story) by making new assumptions about the intentions or capabilities of the enemy. But too many such assumptions render the situation model unreliable. Detecting such unreliability depends crucially on remembering past assumptions (current episodic memory).

More proficient situation assessors may differ from less proficient ones along a variety of dimensions. In terms of knowledge they may differ in the number, automaticity, and power of recognitional templates (structures supporting procedural processing); the richness, organization, and scope of long-term memory structures; the number of cases used to encode exceptions; and the tendency to utilize proactive in addition to predictive or reactive knowledge structures. In terms of metacognitive skills, they may differ in their sensitivity to time and stakes, their propensity to verify the completeness, reliability, or consistency of solutions; and their ability to find the appropriate data source, interpretation, or knowledge for a particular

problem.

Utilization of Findings:

The battlefield situation assessment framework should be of use to a wide spectrum of people concerned with the improvement of situation assessment performance. It clarifies the processes and skills required for successful situation assessment, including high-level principles, knowledge structures, and skills in monitoring and regulating cognition. These processes and skills sometimes diverge significantly from doctrinally prescribed methods or normative approaches. For example, traditional doctrinal guidance to compare qualitatively different courses of action may conflict with the processes of generating, verifying, and modifying a single option that are described within the framework. Similarly, the notion that indicators of enemy intent always have a fixed meaning conflicts with the process of reinterpreting data to fit a coherent, plausible model. On the other hand, the framework addresses aspects of situation assessment for which little if any guidance is currently provided, e.g., determining the contents of the commander's battlespace, or the generation and elaboration of a single effective course of action.

The framework may be used by instructors, designers of C^2 systems, and C^2 analysts and researchers. It can serve as the foundation for the development of a variety of techniques for improving situation assessment performance. Such techniques may include training; the design of supporting materials (such as overlays) or computer-based aids; the improvement of procedures, doctrine, or organizational structure; and personnel selection.

CONTENTS

	Page
Introduction.....	1
Overview of the Framework.....	2
Memory Structures and Knowledge Structures.....	5
Long-Term Memory.....	6
Implicit and Explicit Focus: The Situation Model.....	20
Current Episodic Memory.....	23
The Situation Assessment Product.....	25
Actions, Goals, and Values.....	25
Structural Constraints.....	27
Monitoring and Regulating Cognitive Processes.....	30
Control and Quick Verification.....	31
Full Verification.....	33
Facilitation.....	35
Paths Through the Framework.....	37
Ensuring Plan Completeness and Reliability.....	41
Testing Expectations and Conflict Resolution.....	46
Verifying Assumptions and the Reliability of Data.....	61
Comparing Options and Adjusting Goals.....	63
Modes of Processing.....	72
Situation Assessment Expertise.....	76
Procedural Expert-Novice Differences.....	77
Knowledge-Based Expert-Novice Differences: Long-Term Memory.....	77
Knowledge-Based Expert-Novice Differences: Metacognitive Skill.....	78
Conclusion: Three Approaches to Situation Assessment.....	79
References.....	81

CONTENTS (Continued)

	Page
List of Tables	
1. Modes of processing.....	74

List of Figures	
1. Framework components and relations.....	3
2a. A generic enemy plan structure.....	8
2b. Generic enemy plan structure with associated situation assessment actions.....	10
3. Example of a predictive strategy for dealing with an enemy plan structure.....	11
4. Enemy goal structure.....	13
5. Example of a proactive strategy for dealing with an enemy plan structure.....	14
6. Terrain pattern reflecting a kill zone.....	16
7. Reactive enemy plan structure.....	18
8. Enemy plan execution structure.....	19
9. LTM representation of exceptions within enemy goal structure.....	21
10. Explicit situation model is currently active part of the network.....	22
11. Situation model can be analog.....	23
12. Current episodic memory includes history of situation models.....	24
13. Three types of structural constraints.....	28
14. Monitoring and regulating intuitive behavior.....	32
15. Rapid recognition-primed decision making.....	33
16. Types of uncertainty revealed by verification.....	39
17a. Ways in which solving one problem can lead to another, by adjusting assumptions.....	40
17b. The effect of collecting new data or activating LTM knowledge.....	42
18. Key to symbols in figures.....	43
19a. Progressive deepening: involves filling gaps in plan.....	44
19b. Filling gaps has lead to faulty assumptions.....	45
20a. Testing expectations can lead to conflict.....	47
20b. Resolve conflict by changing assumptions about data.....	48
20c. Search for assumptions underlying data, and acceptance of the simplest way to resolve conflict.....	49
21a. Noticing conflict in initial information about enemy intent.....	51
21b. Conflict resolution gives greatest weight to enemy goals.....	52
21c. Resolving conflict involves constructing the best story.....	55
21d. Resolving conflict by relatively automatic explaining away.....	56
21e. Resolving conflict by searching for unreliable assumptions.....	57
21f. Relatively automatic explaining away based on current episodic memory.....	58
21g. Explaining away can lead to accumulation of too many	

unreliable assumptions.....60

22a. Standard view of terrain generates initial plan.....62

22b. Critiquing view of terrain leads to a new plan.....64

23a. Generation, verification, and modification of initial plan.....66

23b. Detection of conflict and generation of evaluation matrix.....68

23c. Resolve conflict by revising evaluative criteria.....69

23d. Evaluate new plan structure in terms of key criterion.....71

24.Procedural vs. knowledge-based processing.....73

Introduction

Two converging trends -- one theoretical and the other pragmatic -- highlight the growing importance of situation assessment in battlefield command. One set of trends is the increased emphasis in cognitive psychology on pattern recognition rather than explicit analysis, in both problem solving and decision making. For example, empirical studies comparing novices and experts in fields such as physics, chess, and computer programming have supported a view of expertise as the accumulation of direct responses to familiar situations, in contrast to the more analytical, general-purpose strategies adopted by sophisticated novices (e.g., Chase and Simon, 1973; Larkin, 1980). In parallel, research on decision making suggests that experienced decision makers rely on recognitional skills developed over long experience in a domain. Recognition of situations is sometimes associated with direct retrieval of typical responses (Klein, 1993).

The more sophisticated recognition-based theories do not simply equate expertise with the accumulation of a stock of situation templates. First, recent research has emphasized the structure of expert knowledge rather than simply its quantity. For example, many studies suggest that recognition by experts occurs in terms of fundamental domain concepts rather than superficial features of a problem (Chi et al., 1981; Shoenfeld and Herrman, 1982; Weiser and Shertz, 1983; Adelson, 1984; Larkin, 1981). Secondly, there is evidence that experts are more skilled in *metacognition*, i.e., processes that monitor and regulate more basic cognitive processes, like attention, memory, and comprehension (Larkin, 1981; Glaser, 1989; Brown and DeLoache, 1978). Both of these elements -- knowledge structure and metacognition -- increase the flexibility of expert performance and enhance its ability to deal with novel situations.

The second set of trends has to do with evolving flexibility in Army doctrine and practice. An example is the concept of *battlespace* (in the forthcoming FM 100-5). A commander's battlespace is a three-dimensional moving volume that contains anything relevant to his planning or operations. Unlike the traditional *area of interest*, battlespace is not handed down by higher command, but is defined by each commander. It reflects the commander's ability to visualize relevant events at an appropriate level of detail, far enough into the future, and in a large enough volume of space. A second important development in recent doctrine is the decreasing emphasis on enumerating and comparing alternative courses of action. In circumstances of time stress, commanders might use "abbreviated" methods, in which only a single course of action is proposed and assessed. The effectiveness of that course of action will clearly depend on the validity of the commanders' understanding of the situation.

Work is required to integrate these two trends, from cognitive science and Army doctrine. For example, little guidance is currently available as to how Army commanders and their staffs can develop effective representations of the battlespace. Similarly, there is little understanding of how or when the

evaluation of courses of action should be "abbreviated." A systematic investigation of the cognitive dimensions of situation assessment may provide the foundation for improved training, computer-based aids, materials, procedures, doctrine, and perhaps even personnel selection.

This report describes a cognitive framework for battlefield situation assessment by commanders and their staff, based on recent work in cognitive science and on interviews with active-duty command staff. We try to describe the framework, as much as possible, by means of examples taken from the interviews, with members of the command staff at the brigade, division, and corps level. In subsequent reports we will (1) refine, modify, and flesh out details of the framework based on additional interview data, and (2) develop and apply methods for improving key situation assessment skills based on the framework.

Overview of the Framework

Figure 1 depicts the basic components of the battlefield situation assessment framework and the relationships among its parts. At the most general level, the four components of the framework are:

1. memory and knowledge structures;
2. actions, goals, and values;
3. processes for regulating and monitoring cognition; and
4. the real-world environment.

The basic form of the framework and its cyclical character are inspired by Neisser (1976). In his concept of the perceptual cycle, (1) knowledge structures called schemas actively direct (2) attention to and active exploration of the environment. (4) The real-world information generated by that exploration then causes changes in (1) the schemas. These interactions cycle continuously as the observer gains understanding of the actual world. Adams, Tenney and Pew (1991) applied Neisser's perceptual cycle concept to the domain of situation awareness, with some refinements in the characterization of (1) memory structure. Connolly and Wagner (1988) also used Neisser's concept, extending it to include *decision cycles*, in which exploration of the environment causes decision makers to refine (2) their understanding of their goals. In Figure 1, we have incorporated these extensions, and have added (3) the iterative role of metacognition, i.e., monitoring and regulating one's own cognitive processes, in learning both about the world and about one's own goals.

One result of the model in Figure 1 is a somewhat expanded notion of a cycle. As noted in the previous paragraph, Neisser's perceptual cycle comprises only the sequence from knowledge to action to real-world and back to knowledge. Our notion of *cognitive cycle*, however, includes many other, more complex possibilities. As just one example, (1) an initial knowledge structure may be (3) checked by metacognitive processes, (1) modified, and (3) checked again, before leading to (2) an action plan, which is also (3) checked by

metacognition before (3) being implemented in the (4) environment, resulting in (1) new knowledge. We shall refer somewhat loosely to *any* closed loop among the four major components of the model -- for example, knowledge to metacognition to knowledge -- as a *cycle*.

Each major component of the framework will be further broken down into subcomponents:

(1) Memory structures include (a) explicit focus (representing the currently attended part of the situation), (b) implicit focus (containing the full situation model), (c) current episodic memory (containing the history of the current problem), and (d) long-term memory (with both semantic and episodic contents). We will discuss a variety of examples of long-term memory knowledge structures that are used to organize situation assessment information. *Enemy plan structures* organize information about enemy goals, strengths, and opportunities, and describe how they lead to intentions, actions, and consequences. *Enemy goal structures* describe the hierarchical and compensatory relationships among ultimate values, principles, goals, and actions. *Temporal plan execution* structures provide a more detailed description of the temporal durations, precedence relations, and causal contingencies among actions and events. *Enemy planning/C²* structures describe the enemy roles and activities involved in producing, communicating, and implementing plans. *Terrain* structures relate terrain features to expected enemy actions and prescribed friendly actions. Different situation assessors may *frame* problems differently, i.e., bring different knowledge to bear on it. Alternative frames include *proactive*, *predictive*, or *reactive* strategies for predicting enemy intent. Similarly, different assessors may conceptualize enemy plans primarily in terms of goals, terrain, or relative strength.

(2) Action/goal/value structures parallel memory structures, but represent the extent to which events and states of affairs are *desired* rather than believed to exist. Action/goal/value structures include (a) the currently executed or considered part of an action or plan, (b) the current plan, (c) the trajectory image (a sequence of desired situations or major goals for the current problem), and (d) high-level values and principles.

(3) Metacognitive processes shape and guide the retrieval of knowledge from long-term memory and its synthesis in a model and/or plan for the current situation. Metacognition includes components of (a) quick verification (*Is there some reason to think more about my current model or plan, or should I act immediately?*), (b) full verification (*What are the potential problems with the current model or plan?*), and (c) facilitation (*What can I do to improve the current model or plan?*). Quick verification assesses the time available before a decision must be made, the stakes of the decision, and the degree of confidence in the current solution, and determines whether additional processing is required. Full verification tests the current situation model to determine whether it is incomplete, unreliable, or in conflict. Facilitation guides the collection of new data, the activation of additional knowledge from long-term memory, or the revised interpretation of current data to correct any problems that are found.

Once we have defined the basic components of the framework, we can use

them to clarify a broad range of situation assessment phenomena:

- a. structural constraints on processing and the effects of stressors;
- b. different paths through the framework, corresponding to different ways of handling problems that are found in situation models or plans;
- c. different modes of processing that are induced by problem materials or individual differences among assessors; and
- d. characteristics that distinguish expert situation assessors from novices.

These phenomena are the real purpose of the framework. Understanding and dealing with them is where the practical payoff lies, whether in training or in the design of aids and procedures. It will determine how much we can improve a situation assessor's ability to manage stressors, and to select appropriate modes and strategies for processing. Our account of these phenomena should therefore, in an important sense, be regarded as part of the framework itself.

We will try to explain each of these phenomena in terms of characteristics and interactions among the framework's basic components:

a. Structural constraints on processing include (1) the size of explicit focus, (2) time and accuracy in the activation of relevant information from current episodic memory and long-term memory, and (3) effort expended in metacognitive thinking. There are training methods that can mitigate the consequences of each of these constraints: viz., chunking, skilled memory, and overlearning.

b. Paths through the framework involve specific sequences of steps involving the major components of the framework (memory, actions/goals/values, metacognition, and the real-world). We describe a way of analyzing these sequences as combinations of a small number of building blocks. The building blocks are elementary sequences, based on the ways that solving one type of problem can sometimes (but not always) lead to another problem. Among the elementary patterns, for example, that may occur together are the following two: (1) If full verification discovers that data are incomplete, facilitation may fill the gap by adopting assumptions. (2) In a subsequent cycle, full verification may conclude that these assumptions are unreliable, and facilitation may drop them. (Other elementary steps involve adjusting assumptions to resolve conflict; collecting new data to resolve conflict or fill gaps; and activating information in long-term memory to resolve conflict or fill gaps.) Identification of patterns of elementary sequences can play a major role in the design of training or aiding techniques that guard against common situation assessment errors. For example, training might encourage situation assessors to keep track of how many steps have involved assumption adoption, and when the number of assumptions is large, to use full verification to check the reliability of those assumptions. As an aiding example, the success of full verification in detecting unreliable assumptions depends on accurate recall of previous situation assessment steps from current episodic memory; computerized displays might graphically depict previous data and the con-

clusions based upon them.

c. Modes of processing can be distinguished at a general level along two dimensions. They may be *procedural* (involving no metacognitive full verification or facilitation) or *knowledge-based* (involving one or more metacognitive cycles). And they may be *intuitive* (involving knowledge structures acquired by experience) or *analytic* (involving knowledge structures explicitly taught or developed). Four possible modes of processing, then, are: (1) procedural intuitive, (2) procedural analytic, (3) knowledge-based intuitive, and (4) knowledge-based analytic. These modes of processing differ in the demands they place on structural constraints, in the kinds of long-term memory structures they employ, and in the types of metacognitive monitoring and regulation that is most appropriate. Both training and aiding should be designed to encourage the problem-solving approach that is appropriate for the problem at hand.

d. Characteristics that distinguish situation assessment experts from novices may be described in terms of the framework: (a) Experts may differ in the proficiency of procedural processing (e.g., more recognitional templates, more chunking, and more automatic responding); (b) experts may differ in the types of long-term memory knowledge structures that they utilize (e.g., more detailed and more extensive causal models, a larger repertoire of specific cases, and more use of proactive strategies); and (c) experts may differ in metacognitive skills (more explicit attention to time, stakes and confidence; better strategies for finding problems with a model or plan; and better strategies for fixing such problems). These differences between experts and novices help us define the goals of a training program that transfers the skills of more proficient situation assessors to those that are less proficient. It also helps define the functions of a computer aid that guards against errors associated with less proficient processing and encourages the strengths of more proficient processing.

The current state of cognitive science does not permit a definitive catalog of the cognitive components and processes that underlie situation assessment. Applications, however, may benefit from a systematic presentation and incremental advancement, however imperfect. The present framework is still under development, and will be refined and modified as more observational, interview, and experimental data are obtained.

Memory Structures and Knowledge Structures

The left hand pie slice of Figure 1 uses a distinction among four types of memory (A slightly different version of this distinction was discussed by Adams, Tenney and Pew, 1991, and originated with Sanford & Garrod, 1981). The four types of memory vary in the effortfulness with which their contents can be accessed. We distinguish among:

- ! Explicit focus or working memory, i.e., the portion of the situation model currently attended. This is the most activated part of long-term memory.
- ! Implicit focus, i.e., the full situation model. This is a somewhat less activated part of long-term memory, but is readily available for attention.

- ! Episodic memory of the current problem, i.e., the history of situation models and steps in the current context. This is the potential for recalling the preceding series of understandings, plans, assumptions, and choices.

- ! Long-term memory, which contains both general (or "semantic") knowledge and episodic (or case-based) knowledge. (We regard episodic knowledge of the *current situation* as part of a separate memory, since it is usually -- though not always -- more readily accessed than episodes that are more remote in time. It also plays a special integrative role in many situation assessment tasks.)

The distinctions among these four memories are not meant to be absolute, and the boundaries between them can be fuzzy. A helpful metaphor is the connectionist conception of retrieval as *activation*. Thus, explicit focus, implicit focus, current episodic memory, and long-term memory are differentially activated parts of the same system rather than separate compartments (McClelland & Rumelhart, 1986). The knowledge stored in the system can be regarded as a network of interconnected concepts or propositions.

We now consider in turn some of the knowledge structures that characterize each of these four memories.

Long-Term Memory

Long-term memory is relatively permanent knowledge that includes both general relationships and specific past episodes or experiences. General relationships include relatively *analytical* information, such as mathematical rules, as well as relatively *intuitive* information, such as prototypes representing the likely properties of objects or events. Analytical information is typically learned by explicit instruction, or by explicit reasoning using rules that were learned by explicit instruction. Intuitive information is typically acquired by experience. Prototypes, for example, may be built up over repeated experiences with a given type of object or event, during which common or unvarying properties become accentuated, and unique or varying properties cancel out. It may be quite difficult to articulate intuitive knowledge explicitly. (We return to the distinction between analytic and intuitive processing in the section below on "Modes of processing.")

Intuitive information in long-term memory includes not only generalized experiences, or prototypes, but also specific episodes, or exemplars. A key function of long-term episodic memory is to record exceptions to the general rules in semantic memory. Schank (1982) postulated that episodes are stored when expectations based on the generalized experiences fail. These episodes are then tagged by explanations of the failure. The notion that specific memories fill in the gaps in general memories can also be derived from connectionist learning rules. For example, the generalized delta rule or back-propagation learning mechanism predicts that new associations are learned when events are surprising (Rumelhart, Hinton, & Williams, 1986). Thus, the unique parts of episodes will be encoded with specific reference to time and place of occurrence, whereas the non-surprising parts will be merged within a more

general representation, differences of time and place canceling out. Kosslyn and Koenig (1992) cite evidence that different types of neural network, located in different parts of the brain, may be differentially tuned for the storage of semantic versus episodic memory; i.e., some networks may be more readily "surprised," and thus tend to store more unique information.

We will utilize a simple heuristic vocabulary for representing knowledge: as a graph of directed, signed connections among symbolic structures or hypotheses (as in recent work by Pennington and Hastie, 1988); Holyoak, 1991; Pearl, 1988; and many others). A positive connection between two propositions means that the truth of one causes or predisposes the truth of the other; a negative connection means that the truth of one inhibits or predisposes against the truth of the other. In this network metaphor, *schemas* can be interpreted as sets of positively connected units which are likely to be active at the same time (Rumelhart, Smolensky, McClelland, & Hinton, 1986; Schank, 1982). Schemas are thus not regarded as discrete, hermetically sealed packages of knowledge. Membership in a schema is graded rather than all-or-nothing; schemas may come into being gradually as experience modifies connections; different schemas may overlap in their membership; and multiple schemas may be active at the same time.

We will provide examples of three types of long-term memory knowledge structures that are used in battlefield situation assessment:

- ! enemy plan structure
- ! enemy goal structure
- ! enemy plan execution structure.

These are not separate compartments in memory; rather they are all parts of the same interconnected network, with characteristically different but complementary uses. We will also discuss a somewhat less general structure used for visualizing actions and decisions in relation to terrain.

These examples are taken from interview data collected in this and other projects. They are presented to show how situation assessors use their knowledge about enemy goals and values in somewhat different ways, to build situation assessment models that can serve as a basis for developing plans. Each example represents an approach that was regarded as useful by a specific individual to handle a specific problem. Together, however, they might be incorporated into a training program to provide a spectrum of tools that would be useful in situation assessment.

Figure 2a depicts a generic enemy plan structure. It is modeled after Pennington and Hastie's (1988) story model of juror decision making. They propose that jurors construct stories to organize and account for evidence that may be presented in a sequential but nonstory-like order in the course of a trial. The main components of our structure are the three boxes at the top of Figure 2a, representing interests, relative strength, and location. In terms of the traditional METT-T categories (mission, enemy, own troops, terrain and weather, and time), the *interest* box includes the enemy's mission,

high-level motivating values, principles or doctrine, and other high-level enemy goals; the *strength* box corresponds to the enemy's troops and assets in relation to our strength; and the *location* box includes the locations of enemy troops and enemy objectives with respect to one another, to our troops, to the terrain, and to weather, insofar as they affect mobility and combat effectiveness. In courtroom language, the three boxes represent motive, means, and opportunity, respectively. These three elements cause or predispose an enemy intent that satisfies the interests within the constraints of strength and location. Intent, in turn, causes enemy action to realize the enemy interest, and the actions in turn lead to consequences.

We have found evidence that the three principle elements (interests, strength, and location) are regularly considered in building an understanding of enemy plans (although their use may not be perfectly systematic or invariant). The inclusion of these components in a training course or aiding display might improve the situation assessment process. They represent highly general tools for organizing information to explain and predict enemy actions.

The enemy plan structure does not impose any particular order of inferencing. Parts that are activated (or instantiated) first will, in conjunction with other knowledge, lead to activation of other parts. Specific ways of using the structure, i.e., styles of constructing stories, may be associated with its components, just as actions are attached to knowledge structures in schema theory. Three different ways of dealing with enemy plans are depicted in Figure 2b by the bold shadowed boxes:

- ! Proactive: The commander shapes the battlefield, attempting to influence the enemy's intent by altering aspects of interests, strength, or location. In effect the commander causes a story (i.e., enemy plan) of his own choosing to be instantiated.
- ! Predictive: The commander uses knowledge of the enemy interests, strength, and/or location to predict the enemy's intent and actions.
- ! Reactive: The commander infers enemy intent (and possibly interests) after the fact by observing the actions that the enemy executes to achieve the intent or by observing the consequences of those actions.

These strategies are not mutually exclusive. A predictive strategy may employ reactive methods (i.e., observations of enemy actions) to confirm the predictions. A proactive strategy may use predictive methods to decide what actions would produce the desired enemy intent, and may use reactive methods to confirm that the attempt to influence enemy intent was successful.

All three of these strategies support friendly actions designed to exploit or disrupt enemy actions or their consequences. However, proactive, predictive, and reactive strategies will lead to very different outcomes in terms of seizing and maintaining the initiative.

Predictive strategy. Figure 3 is an example of the predictive use of the enemy plan structure. In this example, there are two successive applications of the enemy plan structure, each associated with a prediction. First, from

the fact that the enemy wants to advance and that they have the strength and location to advance, it is a simple expectation that the enemy will lay in or bring up fuel supplies. This prediction triggers an action by friendly forces: destruction of the enemy POL (petroleum, oil, and lubricants) dump by the Air Force. The success of this action in reducing enemy fuel supplies is confirmed by observing the subsequent actions of the enemy. In particular, the enemy stops running trucks at night.

The second application of the predictive enemy plan structure is more interesting. Since the enemy is low on fuel, it is now inferred that an enemy interest is to obtain fuel. Combining this with the observation that friendly POL is accessible on the road in front of them (location) and that the enemy have adequate strength to attack our POL, the commander now expects that the enemy intends to attack our POL and seize our fuel supplies. As a result of this prediction friendly forces defensively strengthen themselves in the area of the POL depot.

Plan structures like the one in Figure 3 can play a role in *procedural* processing, in which a situation is directly recognized and an appropriate action retrieved. If the situation assessor has experienced incidents similar to this one in the past, a knowledge structure like Figure 3 might already exist, ready to be activated. For example, if the features of the present situation (e.g., the intent to advance, strength, and location) match the plan structure at the top of Figure 3, then expectations regarding the enemy's laying in fuel supplies are automatically activated, along with the friendly option of finding and destroying the enemy POL dump.

On the other hand, plan structures of this sort can also play a role in *knowledge-based* processing, in which appropriate representations and/or responses do not pre-exist, but must be constructed from existing knowledge. In these less familiar situations, the assessor may *construct* a relevant plan structure from the generic plan structure in Figure 2 together with other knowledge. (We return to the distinction between procedural and knowledge-based processing in the section, "Modes of processing.")

Figure 4 depicts one kind of knowledge structure that may help to fill in the plan structure at the bottom of Figure 3. It depicts in some detail an enemy goal hierarchy, starting with a major goal at the top and moving down to more specific actions at the bottom. Note that this particular structure pertains to one expert's assessment of one enemy's goals in one type of situation. But it may be illustrative of a class of representations that is widely useful.

Figure 4 contains three types of elements: major goals or objectives, principles, and actions. The top level of this structure represents the *trajectory image* of Beach (1993). The enemy has a series of major goals or objectives that he wishes to accomplish, the first of which is achieving a successful attack (penetration) in a particular region within a particular time window. Subsequent goals might include seizing a particular terrain objective (e.g., a city) within a later time window, etc. The expert from whom this structure was elicited further organized his understanding of enemy actions in terms of two more general, higher-level principles: increase the

capabilities of the enemy's forces and reduce the capabilities of the enemy's opposition. Each of these higher-level principles was subdivided into principles that focus on location, i.e., concentrate enemy forces or disperse opposing forces, and strength, i.e., increase the size of enemy forces or reduce the size of the opposing forces. Such a hierarchical structure, in which actions are organized by objectives and high-level principles for achieving those objectives, can be quite useful in predicting the kinds of plans that the enemy will develop in different situations. For example, if the enemy is unable to increase the size of its own forces in the battle area, it may compensate by drawing off some of the opposing forces by means of a diversionary attack. In the present example, if the enemy is unable to use its own fuel supplies, it may attempt to seize supplies from us. (The latter might be an instance of another more general enemy principle: If you fail to secure your own supplies, seize supplies from the other side.)

Figure 10 depicts the goal structure in Figure 4 side by side with the plan structure of Figure 3. Numbered arcs in Figure 10 show how activation in the plan structure might be mediated by connections in the goal structure. In particular, the goal structure might enable the assessor to make the connection between destruction of the enemy's POL dump and the enemy intent to attack our POL. For example, (1) the expectation that the enemy will lay in or bring up its fuel supplies (in the top plan structure) might activate the node representing *create own POL depot* in the goal structure. (2) When the enemy POL dump is destroyed, the higher level goal *POL* is activated. (3) This in turn causes activation of the interest node in the lower plan structure (obtaining fuel is now an enemy priority). A governing principle (seizing other's supplies if you lose your own) might also be activated in the goal structure (not shown in Figures 4 and 10). This principle plus the *POL* goal node lead to activation of the alternative subgoal, *seize enemy POL*, in the goal structure. (4) That subgoal matches the criteria associated with enemy intent in the generic plan structure (Figure 2); and so it activates the intent node in the lower plan structure. The assessor may now gather information to fill in or verify other nodes in the plan structure, e.g., he may examine enemy capabilities and position in relation to our POL depot.

Proactive strategy. Figure 5 is an example of the proactive use of an enemy plan structure. The top half of Figure 5 represents the *friendly* command staff plan structure. The friendly mission (i.e., interest) in this scenario was to defend at a particular phase line. The situation assessor considered the relative strengths of the two sides: A force ratio of 3 to 1 favoring the enemy yielded a chance of successful defense of only 50 percent. Most important of all, the situation assessor considered location in an active rather than a passive fashion. He actively looked for terrain possessing certain features, which he defined as a kill zone: an open area which is accessible only through restricted terrain, for which there is an early trigger point (indicating whether or not the enemy is headed to the kill zone), and to which he can maneuver with limited visibility by the enemy. In this scenario, he found such a zone in the north. These three factors, a defensive mission, less than advantageous relative strength, and the discovery of an appropriate kill zone, lead to an intent: to induce the enemy into the kill zone.

But how can the enemy be induced into the kill zone? The key to developing

an appropriate friendly action is to use the *enemy* plan structure and create an equation between what the enemy wants and what the friendlies want the enemy to do. The enemy plan structure in the bottom half of Figure 5 can thus be used to fill gaps in the friendly plan structure in the top half of Figure 5. We say that the enemy plan structure is being used proactively, because the assessor attempts to change the perception of relative strength in the enemy plan structure, in order to produce a desired enemy intent.

Like the predictive plan structure considered earlier, the proactive structure may be utilized in either a procedural or a knowledge-based process. For an assessor who is not familiar with aspects of this situation, the structures in Figure 5 may not pre-exist, but may be constructed, at least in part, through an incremental knowledge-based process. In this case, the assessor *derives* some of the required properties of a kill zone, and/or the tactics for inducing the enemy into it, from a causal understanding of the enemy and the terrain. Such a process is not linear or unidirectional. It is a matter of constructing a friendly plan and a representation of the enemy plan in a parallel and mutually constrained manner.

For example, based on his own defensive mission and poor force ratio, an assessor might initially look for a kill zone defined by only two properties (an open area with narrow access). Noticing that the open area in the north satisfies these conditions, he might form the intention of inducing the enemy into it. To figure out how, he activates a representation of the enemy plan (the bottom section of Figure 5). The situation assessor begins with a high level Soviet principle or doctrine that determines where to go: avoid opposing combat power. In this situation, that principle translates into a high-level enemy goal: to go where there are the least tank killing systems. In terms of location, the Soviets have equal capability of going north or south. Thus, if the Soviets perceive the least opposing combat strength in the north, this plan structure predicts that they will develop and implement an intent to go north. The friendly intent then is determined: to place the fewest tank killing systems where they want the enemy to go, i.e., in the northern open area. In other words, friendlies make themselves look weak in the north. This translates into the specific action of putting a light brigade in the north and a heavy brigade in the south.

Now the assessor may verify this plan by mentally simulating it. Supporting this mental simulation might be a representation of events in relation to terrain and time, as shown in Figure 6. Numbers in Figures 5 and 6 correspond to the temporal order of events in the mental simulation. The mental simulation helps the assessor notice important failures and gaps in his current plan. First, (items 1 and 2 in Figures 5 and 6), in order to set the trap, the enemy must *believe* that friendlies are weak in the north. They can only know this through their own reconnaissance. This will not happen if friendlies kill all the enemy reconnaissance. This leads to the friendly action of not killing all the enemy reconnaissance. Second, (items 3 and 4) the assessor must get the heavy brigade north in time to meet the enemy there. How will he know when to move it? This implies an additional requirement for a kill zone: that it possess an early trigger point, at which the enemy must commit itself to going north. It also implies the friendly action of assigning reconnaissance to the trigger point. Thirdly, (items 5 and 6) when the

assessor does move the heavy brigade north, if the enemy detects the movement, they may not enter the kill zone. This implies the requirement that the maneuver area be shielded from enemy observation. The assessor can now verify that the terrain initially selected as the kill zone satisfies the two additional constraints (early trigger point, unobservability) suggested by knowledge-based processing. (Mental simulation to verify plans will be discussed at greater length in the sections "Monitoring and regulating cognitive processes" and following.)

If the assessor is familiar with this type of situation, many of these steps can occur within relatively automatic, procedural processing. In this case, structures like Figure 5 and 6 may largely pre-exist and be activated by features of the current problem. For example, the defensive mission, poor force ratio, and knowledge that the Soviets are the enemy might activate the search for a kill zone with pre-defined properties. Rather than being dynamically constructed, Figure 6 might exist as a ready made template (tailored for a Soviet-style enemy) that the situation assessor applies to the terrain. This structure consists of four key components: an open area, narrow access, an early trigger point, and the possibility of maneuvering forces to the kill zone out of sight of the enemy. (Notice that all he really has to look for, in this compiled version, is an open area in a particular relation to the enemy and to mountains. Mountains serve a triple purpose: narrowing access, forcing an early decision to go north or south (a trigger point), and blocking visibility.) Finding a suitable kill zone then directly triggers the actions of placing the friendly light brigade in the kill zone (the north), the heavy brigade in an adjacent area (the south), and friendly reconnaissance at the trigger point.

Proactive uses of enemy plan structures can take other forms. In this example, the friendly strategy was to influence enemy intent by manipulating the enemy's perception of relative strength. In the section, "Ensuring model or plan completeness," we will describe an example of a proactive strategy in which enemy intent was influenced by manipulating enemy interests, by disrupting communication of goals from higher command (Figure 19a). In that example, construction of an adequate proactive plan is supported by a model of enemy planning and C² activities.

Reactive strategy. Figure 7 is an example of a reactive use of the enemy plan structure. In this case, enemy intent is neither being influenced nor predicted from higher-level goals and objectives; rather, intent is inferred from actions that are already underway to implement the intent. (For this reason, boxes denoting interests, strength and location may be less relevant and are not shown.) In particular the time and location of enemy actions are often used to calculate the likely time and place of an attack. This calculation is, of course, a prediction; but since it is based on overt enemy movements rather than pre-existing interest, strength, and location factors, we choose to regard it as a reactive strategy. In the example of Figure 7, the conclusion serves to motivate friendly decisions to commit reserves in a particular place and time to prevent a successful enemy penetration.

Like predictive and proactive structures, the reactive plan structure can be used in either procedural or knowledge-based processing. The procedural use

of these structures may be relatively analytic or relatively intuitive. The analytic case is little more than the familiar use of memorized "indicators" of time and place of attack. The intuitive case involves sensing a pattern of activity that has been associated with time and place of attack.

Figure 8 shows a knowledge structure that can support reactive knowledge-based processing. It is an enemy plan execution structure and represents partially constrained precedence relations among enemy actions. For example, this diagram indicates that an increase in reconnaissance activity facilitates the success of a large number of other activities. Similarly, an increase in supply activity facilitates a large number of other activities. Assigning higher echelon artillery to the front echelon attacking unit is a prerequisite for moving that artillery up, emplacing it, and beginning an artillery barrage. On the other hand, assigning higher echelon artillery is not a

necessary precursor for creating a diversionary attack, for cross-attaching tanks, moving up backup units, preparing air defense, moving up engineer assets, etc. Precedence relations of this sort constrain but do not fully dictate the order in which various enemy actions would be expected to be observed.

Figure 8 can be used to infer future observations from present or past observations. For example, if engineer assets are observed being moved up, the situation assessor can conclude that the engineer assets will subsequently be emplaced and used to remove obstacles. If he has also observed units massing in an offensive formation, he can conclude that the sector will be narrowed and that the attack will subsequently take place. Conversely, given a conclusion of the time and place of attack, the situation assessor might work

backward: for example, concluding that radio silence and an artillery barrage will occur at particular times and places. Thus, conclusions from such a structure can be used both to construct, and to confirm or disconfirm, hypotheses about enemy intent.

In addition to the partial constraints among events represented by Figure 8, continuing experience may lead to direct links between events and time to attack (represented by the numbers along the top of the figure). Such associations might be weaker than the precedence relations among the events. But they may afford another way in which this structure could be used to predict and verify enemy intent.

Situation assessment strategies and frames. In sum, we have described three different strategies for situation assessment: proactive (influencing the enemy's intent by affecting his interests, strength, or location), predictive (using knowledge of the enemy's interests, strength, and/or location to predict his intent), and reactive (inferring the enemy's intentions after observing actions he has already executed in order to realize his intent). Each of these strategies can be used in a direct, recognitional way, based on pre-existing knowledge of the relevant factors and their relationships. But each of them can also play a role in knowledge-based processing. In the latter case, they each draw on other knowledge to support the construction of an enemy plan representation. Underlying our example of the proactive strategy was a causal terrain representation of a kill zone. Underlying our example of the predictive strategy was a hierarchical model of enemy goals, values, and actions. Underlying our example of the reactive strategy was a temporal precedence model of enemy plan execution.

A *frame* may be defined as the portion of a decision maker's knowledge that is brought to bear on a particular problem (Beach, 1990). Thus, our discussion of long-term memory knowledge structures has illustrated a variety of different frames. Different situation assessors may frame or conceptualize the same situation differently: e.g., proactively, predictively, or reactively, and in terms of enemy goals, terrain, or temporal precedence. Moreover, the same assessor may bring different frames to bear during different phases of the assessment process.

Exceptions and episodic memory. Referring back to the memory structure portion of Figure 1, we see that long-term memory can store general or semantic structures. It can also, however, store specific episodes. A key function of long-term episodic memory is to record exceptions to the general rules in semantic memory. Figure 9 represents an example from the goal structure of Figure 4. The solid arrows represent the normal or semantic relationships within this knowledge structure. We've added an arrow with a negative sign to indicate that the goal of surprise conflicts with the goal of weakening the enemy by means of an artillery barrage. Emplacing the artillery within range of the opposing force is necessary for the artillery barrage, but it may tip off the opposing force as to where and when the attack will take place. The dashed lines and boxes represent exceptions to the general rule that the artillery will be emplaced in the region of the attack. The artillery might be placed in a different but nearby region in order to enhance surprise. The goal of an artillery barrage may be achieved as well, either because the artillery

is longer range than expected or because the artillery can be moved rapidly at the last minute. Episodes in which artillery was not located in the region of attack would be remembered and tagged with the explanation that applies.

The representation of exceptions within long-term memory plays a key role in mental simulation when there is uncertainty about what events will occur, as we shall discuss later in the section on "Testing expectations and conflict resolution." The appropriate use of exceptions, e.g., to explain unexpected data, is a prime candidate for training and decision aiding.

Implicit and Explicit Focus: The Situation Model

The *explicit* situation model at any given time is the activated portion or portions of the assessor's knowledge structures. This activated knowledge, however, may not include the entire situation model. The situation model at any given time also includes implicit, or less activated, portions of the same knowledge structures, such as those networks illustrated in Figures 3 through 9. These portions may themselves become activated from time to time, either alone or in combination with other portions. Figure 10 illustrates this concept by showing how portions of the predictive plan structure in Figure 3 and the supporting goal structure in Figure 4 may be coactive at a particular time, constituting the explicit situation model, while the remainder of these two structures, which are implicit, define the total situation model. As we have already seen, it also illustrates how the two structures are linked by

patterns of activation, and thus how each structure can mediate connections in the other. For example, the circled portions of the plan structures on the left of Figure 10 are linked by their relationship to the enemy's need for POL in the goal structure on the right.

A situation model of the kind illustrated in Figure 10 is basically propositional. It consists of a set of symbolic structures all of which are activated at the same time. Another type of situation model is analog or iconic. We have already seen an example of such a model in Figure 6, the template representing the terrain features for a kill zone. According to Johnson-Laird (1983) *mental model* representations are isomorphic to the represented state of affairs. As a result of such isomorphism new relationships can simply be read off the representation without the need for an exhaustive listing of represented facts, or for highly abstract rules of inference. For example, in Figure 6 the distance which the enemy will travel from the trigger point to the open area can be directly compared to the distance that the heavy force must travel from the north up to the open area.

According to Johnson-Laird (1983), Kosslyn and Koenig (1992), and others,

images or iconic representations are constructed from underlying permanent knowledge structures. Analog situation models need not represent only spatial relationships. For example, an image or mental model like Figure 11 might be constructed by repeatedly activating the nodes in the temporal plan execution structure in Figure 8. From this representation the situation assessor can directly see that enemy forces should be in position at about the same time that opposing forces respond to an enemy diversion. He no longer needs to follow the indirect path forward to time of attack and then backwards to response to diversion.

Figure 11

Current Episodic Memory

Episodic memory is the record of past autobiographical occurrences. As we noted above, incidents that involve exceptions to a general rule may be especially well remembered. Current episodic memory is part of long-term episodic memory that we single out because it is somewhat more easily activated, and because it plays a special role in problem-solving. Current episodic

memory is a record of the situation models or plans that have been sequentially activated in the course of *the present* problem. Figure 12 illustrates this concept in the context of the predictive enemy plan structure of Figure 3. The circled part of the lower half of Figure 12 shows "today's" situation model as shown in Figure 10. The circled part of the upper half of Figure 12 shows the situation model that was activated "yesterday." Current episodic memory keeps track of the on-going intentions and activities of the situation assessor, from adoption of a goal to its success or failure.

Retention of a record of past situation models and plans is crucial to many problem-solving and decision-making tasks within situation assessment. For example, previously considered and rejected action options may be incorporated into a current plan as contingencies or branches, to enable the plan to handle unexpected situations (Fallesen, 1993). As another example, recall of previous assumptions is crucial in deciding whether or not to explain away a new piece of conflicting data. Unfortunately, previous plans and plan options may be forgotten in the course of an evolving situation, and decision makers often forget the assumptions they have already made in building a situation model. We will return to some of these issues in the section on "Testing expectations and conflict resolution."

The Situation Assessment Product

The situation model as we have defined it is the *product* of situation assessment only in a very narrow sense. It is the moment-by-moment crystallization of the assessor's understanding. The true product of situation assessment is both more extensive and more externalized. Although it may never be made completely explicit, it is reflected in the commander's estimate, in his concept of the operation, in briefings given by the commander or his staff, and in the way that the commander and staff communicate with one another. A shared or overlapping model of the situation may be a prerequisite for successful implementation of the commander's intent by subordinates and for successful coordination within a C² organization (Kahan, Worley, & Stasz, 1989). The situation model in this wider sense is not all present to mind, even implicitly, at one time. Its articulation requires repeated cycles through the situation assessment framework (Figure 1), in which previous situation models are retrieved and activated in turn, key components are extracted and combined with one another, and finally, transformed into a relevant external format. The product of this process is the external situation model, which reflects an integrated multi-level representation over a large area of space and time. It consists of the following components (e.g., Endsley, 1988):

- ! classification of objects, i.e., individuals, platforms, units, and organizations; terrain and weather features
- ! integration of objects into patterns, i.e., specification of activities, purposes and histories; and
- ! projection of patterns into the future, i.e., specifying implications for one's own goals and plans.

Actions, Goals, and Values

The right-hand pie slice of Figure 1 represents the decision maker's *evaluations* of events and actions. In the left-hand pie slice of Figure 1, activation of an event has to do with *belief*; it means that the event is occurring, or else that it is being inferred, assumed, or predicted. In the right-hand pie slice, by contrast, activation of that same event means that it is currently valued, intended, being considered for implementation, or being executed. The four types of evaluative knowledge correspond to different degrees of relevance to the current situation.

Long-term evaluations. This represents relatively permanent knowledge about what the decision maker regards as desirable, important, and worth pursuing. Like Beach's (1993) *value image*, it includes the highest level bedrock values (e.g., defeat the enemy with the least loss of life on the friendly side) and principles (e.g., go where there is the least opposing combat strength, as illustrated in Figure 5; influence enemy intentions before the fact whenever possible, as illustrated in Figure 2b, proactive strategy). Its contents may vary in generality; it may include specific goals or even types of actions, when these are pursued or valued for their own sakes rather than as means to an end. Evaluations can be represented quantitatively as degrees of preference, or qualitatively as binary states of affairs or constraints. Long-term evaluations are typically "semantic," i.e., they refer to the value of general types of events or states of affairs. But in rare instances, they can be episodic, i.e., referring to a specific, uniquely valued event in the past or future; e.g., to avoid one's first defeat in battle.

Current episodic goals. These are the decision maker's goals in the current problem. Goals are concrete ways of realizing high-level values and principles. Like Beach's (1993) *trajectory image*, such goals stretch back to the beginning of the problem and project forward in time to its conclusion. It may include, for example, a series of terrain or engagement objectives stretching out in time (e.g., "successful penetration of enemy defenses at phase line x by day d, arrival at objective y by day d+n..." at the top level of Figure 4). These goals extend beyond the immediate situation, but provide its evaluative context. Current episodic goals specify the decision maker's overall set of intentions, the events he or she would like to bring about, and which give a larger meaning to his current actions (e.g., emplacing a POL depot) than their specific intents. Current episodic goals are used in turn to generate plans for the achievement of those goals.

Current plan. This is the detailed set of actions and action contingencies that the decision maker has adopted in the current situation. Like Beach's (1993) *strategic image*, it includes the specific actions (e.g., "move up follow-up forces," "emplace artillery," look for a kill zone, etc.) undertaken to realize goals. As we descend to more specific levels, temporal relationships become increasingly important. Thus, values and principles are relatively time independent; goals may be arranged in a relatively simple temporal sequence or trajectory; and actions require a far more detailed representation of durations, temporal constraints, and contingencies. The current plan may be represented by a plan structure as shown in Figure 2a, and - at a greater

level of detail - by a plan execution structure as shown in Figure 8. Adoption of plans depends far more than values or goals on the assessor's representation of the specific current situation. An action may be adopted because it is linked causally to a goal in the situation model.

Active part of plan. This is the portion of the current plan that is the immediate focus of evaluation. Such evaluation may occur prior to implementation as part of the decision making process, or during implementation by monitoring an on-going action for its success in achieving goals.

We have noted how the representation of our own values, goals, and actions can take exactly the same form as the representation of enemy goal structures, plan structures, and plan execution structures. Moreover, values, goals, and actions can be arranged hierarchically in relation to one another in a structure like Figure 4. Processing in such a structure can be either top down or bottom up (Beach, 1990). Actions and plans may be generated and evaluated based on goals, and modified or rejected if they fail to achieve them. Similarly, goals may be generated and evaluated based on values. On the other hand, from a bottom-up perspective, goals may be revised if no actions can be found to achieve them. Even high-level values might be revised if they are not achievable by realistic goals or actions. We will describe an example of revising goals in the section, "Comparing options and modifying goals."

Decision making is sometimes depicted as a sequential process in which situation assessment is completed before course-of-action generation and evaluation begin. But Figure 1 does not imply that knowledge about values, goals, and actions is separate from knowledge about the world or the situation. Indeed, according to recognition theories, goals and appropriate actions may be directly associated with the knowledge structures that are used to understand situations. For example, Figure 3 showed how prediction of enemy goals (the need for POL) leads, in the context of the higher-level values associated with a defensive mission, to adoption of the intent to block those goals and the action of destroying the enemy POL depot. Moreover, decision-related knowledge structures, like scripts (Schank and Abelson, 1977), embody considerable knowledge about the environment, as well as about actions and their consequences. Figure 1 represents procedural processing in terms of direct links between the activation of situation models in the left-hand pie slice and the activation of values, goals, or actions in the right-hand pie slice.

In knowledge-based processing, course-of-action selection and situation assessment are even more inextricably intertwined. Initial situation understanding may lead to the activation of high-level values or goals. Fleshing out the details of a plan of action, however, may require additional elaboration of the situation model. Figure 5 illustrated how the goal of luring the enemy into a kill zone led to activation and elaboration of a situation model representing enemy doctrine and beliefs. Situation assessment in this case is the means by which the planner converts his goals (e.g., lure the enemy into the kill zone) into a specific plan (look weak in the area of the kill zone, do not kill all the enemy reconnaissance, assign friendly reconnaissance to the trigger point, etc.). As the situation model is fleshed out, components of the plan are activated, and the overall plan design takes shape. Plans are

constructed during, not after, situation assessment. Once a plan is formulated, its adequacy is assessed metacognitively by predicting its outcomes in the situation model (Does it achieve relevant goals?). Goals are assessed in the same way, by longer-range predictions in situation models (Do the goals achieve high-level values?) Conversely, the adequacy of the situation model is assessed by reference to its ability to generate, constrain, and evaluate plans and goals. Later, in the section on "Ensuring model or plan completeness," we will show again how problems with a plan lead to the elaboration of situation models.

While actions, goals, and values are not a separate compartment of knowledge, they do reflect a qualitatively different way of viewing knowledge. The situation assessor needs a representation of how much possible states of affairs are *valued*, in addition to the representation of how much they are *believed*. Values or preferences are importantly different from strengths of belief, and their propagation or pattern of activation through a network will be different from the propagation of strengths of beliefs. As we have seen, each will influence the other. Moreover, once a plan is adopted, if it is expected to be successful, the actions and outcomes of the plan will be activated as expected events in the left-hand pie slice. Moreover, the *consideration* of a plan is an actual event that may be recorded in the left-hand pie slice, as part of the history of problem solving activity in current episodic memory, whether or not it is adopted, i.e., activated on the right side.

Structural Constraints

At the most general level, situation assessment success is constrained by three factors: (1) the inherent unpredictability of real-world events even with all available knowledge; (2) failure to possess potentially available knowledge; and (3) flaws in the use of the knowledge that is possessed. We will discuss (1) and (2) in the section, "Monitoring and regulating cognitive processes," where we focus on how metacognition grapples with uncertainty through adoption of assumptions, activation of knowledge, and data collection. In this section, we focus on (3): built-in, or structural, limitations on humans as information-processing systems.

Three types of structural constraints affect the operation of the situation assessment model (Figure 1). These are shown in Figure 13:

- ! limitations on attention,
- ! limitations on the activation (or retrieval) of knowledge, and
- ! cognitive effort.

Each of these structural constraints is associated with methods for mitigating its effects to at least some degree.

Constraints on attention reflect the size of explicit focus (or working memory). These constraints can be mitigated in two ways: (a) Different modalities of representation, e.g., auditory and visual, interfere less with one another than representations of the same type, e.g., all information

represented visually (Allport, Antonis, & Reynolds, 1972) (b) A large quantity of information may be encoded into a single relational pattern, or "chunk," and thus may be represented by a single token in active memory (Miller, 1967; Newell & Rosenbloom, 1981). Despite these potential mitigations, constraints on the total content of attention are significant factors in performance (Adams, Tenney, & Pew, 1991).

The second structural constraint involves the process of activating information in long-term memory. Such activation may involve errors, i.e., both misses and false alarms. Even when accurate, such activation may take time. This constraint interacts with the first. If active memory were infinitely large, there would be no need for activation from long-term memory. Conversely, if activation from long term memory were instantaneous and error free, constraints on the size of working memory would be irrelevant.

This constraint, like the constraint on the size of explicit focus, can be mitigated. Ericsson and Polson (1988; Chase and Ericsson, 1981) describe evidence for a theory of *skilled memory*, by means of which information can be rapidly and accurately retrieved from long-term memory. Skilled memory involves the use of existing long-term memory structures to encode new data. A key point is that retrieval cues are associated with the new material in the encoding stage. With practice in a specific domain, encoding of new information in that domain can be accomplished very rapidly. As a result, the new material can be rapidly accessed in situations that match the retrieval cues even after long periods of time. We suspect that aspects of skilled memory characterize expertise in a variety of domains. Nonetheless, constraints on long-term memory access are still critical in non-routine tasks for which appropriate retrieval cues have not been prepared in advance.

A possible third constraint involves cognitive effort. More effort is required by higher level executive processes and less effort by more automatic processes of perception and pattern recognition (Shiffrin & Schneider, 1977). Effortful processes may be common in monitoring and controlling the activation of information in long-term memory and its manipulation in explicit focus.

This constraint can be mitigated by extensive practice with consistent stimulus-response mappings (Shiffrin & Schneider, 1977). Practice in a task makes it more procedural or "automatic": Long reasoning sequences, which involve repeated activation of new information in explicit focus, are replaced by direct stimulus-response linkages (Anderson, 1982).

A stressor is any event in the actual present environment which affects any of these three structural constraints. First, a stressor almost always takes up space in active memory. For example, noise seizes attention involuntarily; a high risk situation seizes attention because of the stakes or goals that are affected. Second, a stressor may increase access errors or access time from long-term memory. For example, difficult tasks may themselves be stressors because they impose the need to activate relatively inaccessible material in long-term memory; secondary tasks or stimuli may function as stressors because they lead to cross-talk between their own cognitive activity and activity associated with the primary task. Third, a stressor may consume mental effort by adding executive tasks. For example, metacognitive processes may monitor and regulate the activation and interpretation of information from

long-term memory; alternatively, reasoning may be required to construct a plan to eliminate the stressor or to avoid situations in which the stressor exists.

All three types of structural constraints can degrade situation assessment; as a result, the significance of events may be missed and the situation misunderstood. In some cases, appropriate knowledge structures exist, but they may not be currently active because of constraints on the size of explicit focus, or because they have been displaced from explicit or implicit focus by stressors. If the appropriate knowledge structure is not currently active in explicit or implicit focus when a critical event occurs, and if relevant retrieval cues have not been associated with the knowledge structure, the significance of the critical event may be misunderstood (Adams, Tenney, & Pew, 1991). Successful situation assessment will be limited by the time and accuracy of activating the displaced material and by the effort demanded by metacognitive control processes. In other cases, appropriate knowledge structures do not pre-exist at all; appropriate structures must be activated and combined to construct a situation model or plan. In this case, constraints due to both access and effort are even more severe.

Monitoring and Regulating Cognitive Processes

The third major component of the situation assessment framework is metacognition: the monitoring and regulating of one's own cognitive processes. Metacognition primarily supports knowledge-based, as opposed to procedural, processing. Metacognitive skills may be thought of as a set of techniques for dealing with the structural constraints that limit the effective application of knowledge.

As we saw in the section on "Actions, goals, and values," in some cases there is a direct link between the situation, knowledge structures that are strongly activated in that situation, and goals and actions that are associated with those knowledge structures. Such cases have been called recognition-primed decision making by Klein (1993), rule-based behavior by Rasmussen (1993), and procedural knowledge by Anderson (1982). In other cases, however, previous experience does not provide a ready-made response or problem solution. As a result, *knowledge-based processing* is far more affected by each of the structural constraints discussed in the previous section. Knowledge-based processing is characterized by the following features:

- ! The knowledge structures that must be accessed may exceed the capacity of implicit or explicit focus, and thus require repeated cycles of activating long-term knowledge, extracting or abstracting relevant components of information, and integrating those components into an evolving situation model or plan.
- ! An adequate model or plan may require the activation of information that is only indirectly or tenuously linked to the present situation; it may also require the integration of knowledge structures that are not strongly linked to one another. Knowledge-based processing is thus constrained by limitations on the accuracy and timeliness of long-term memory activation.
- ! Knowledge-based processing is more effortful and time consuming than

procedural processing, since it typically involves metacognitive activity to support the activation and manipulation of long-term memory knowledge.

One function of metacognition in the R/M framework is to determine *when and if* knowledge-based processing is required or justified, because of available time, the costs of errors, and inadequacies in the procedural approach.

Two additional functions of metacognition provide support for knowledge-based processing once it is undertaken. By definition, knowledge-based processing is indirect. The required knowledge is not all present in explicit or implicit focus, nor is it linked to the situation by direct retrieval cues (as in skilled memory). The required knowledge must be searched for, activated, and perhaps abstracted and combined with other information in order to produce a solution. Knowledge-based processing can thus benefit from: (1) control over the way in which knowledge is searched for and manipulated, as opposed to random "free association;" and (2) verification of the solutions that result from this process. These processes reflect metacognitive skill.

Figure 14 shows how metacognitive skills support knowledge-based behavior. It expands the component of Figure 1 labeled *monitoring and regulating*. As already noted, the three major functions of metacognition in this framework are

- a. to determine whether more extensive knowledge-based processing is justified,
- b. to verify the adequacy of the current situation model and/or plan, and
- c. to facilitate improvements in the current situation model or plan.

Since the primary purpose of metacognition in this context is to support and extend recognitional behavior, we refer to Figure 14 as the Recognition/Metacognition model. We now consider the three major functions of metacognition in more detail.

Control and Quick Verification

This step asks three questions: (1) How much time do I have before it is necessary to commit to a decision? (2) How high are the stakes of an error? And (3) is there any reason to doubt my initial understanding or plan for this situation? In quick verification, reasons for doubt are straightforward and do not require extensive activation of long-term memory; e.g., the situation is relatively unfamiliar or atypical in some way, or the assessor is already aware of a problem with the model or plan, such as incompleteness, unreliable data or assumptions, or conflicting data or opinions. If the answer to any one of the three questions is no (i.e., no time is available, or the stakes of making an error are low, or the situation is highly familiar and typical and no problems have been identified), then no further metacognitive processing is required. Procedural processing is adequate or necessary. Figure 3 provides an example, in which recognition of a situation involving the goal of enemy advance directly activates the enemy goal of laying in POL and the friendly

intent to destroy it.

According to Klein (1993), rapid recognition-primed decision making is expected under conditions of high time pressure. According to Connolly and Wagner (1988), it may occur when there is low cost of an error. According to both Klein and Rasmussen, it is expected in highly familiar situations, or from decision makers with high levels of expertise. These three characteristics correspond to the three questions posed by the control and quick verification step. If any of these conditions is satisfied, no further verification or facilitation takes place (Figure 15). The initial model or option is accepted. An implication of this observation is that even in the most rapid recognition-primed process, quick verification must be included. Quick verification may operate in a management by exception mode, working in parallel with direct procedural processing, but inhibiting the execution of the response if problems are found. The quick verification step is, as the name implies, extremely rapid and virtually automatic.

Skilled quick verification may contribute to differences in performance between experienced and inexperienced decision makers. One such difference concerns the timing of decisions. For example, in research on ship-based anti-air engagement decisions (Cohen, 1993), critical incidents involving engagement decisions against approaching targets of unknown identity or intent were analyzed. More experienced officers tended to wait longer before deciding to engage than less experienced officers. The more experienced officers were more likely than the inexperienced officers to explicitly estimate the amount of time available for decision making (e.g., before the target was likely to attack, or before own ship weapons would be unable to counterattack). In a study of the commercial air context, Orasanu (1990) found similar differences in the performance of cockpit crews flying a simulated 737 scenario. There was a tendency for better performing air crews to take more time for decision

making when such time was available, but when time was associated with unnecessary risk, better performing crews acted sooner.

Full Verification

If quick verification fails (because time is available, the stakes are high, and there is reason to doubt the initial solution), metacognitive processing continues. The next step depends on the reasons for doubting the model or plan. If the situation is atypical or unfamiliar, but no specific problem has been definitely identified, the process of full verification begins to look for specific problems. Full verification consists of one or more of three highly intertwined component processes.

Incompleteness. The first process tests for incompleteness or gaps in the

model or plan, and for data which the model or plan have not taken into account. The methods used in the test for completeness may include, for example, mental simulation of future events based on the current model to determine if it accounts for all the observed data, or to determine if it predicts future events at the required level of detail (e.g., to support generation of an adequate plan). A plan might be mentally simulated to make sure that it achieves all relevant goals. Other tests for incompleteness include use of a checklist, a template, or standard operating procedure which details the model or plan components, or the steps that must be followed in producing a model or plan.

Knowledge-based and procedural processing differ in the relationship between situation assessment and courses of action. In knowledge-based processing, the initial situation model is not directly associated with an acceptable course of action. Nevertheless, the situation model is usually associated with *constraints* on possible courses of action. As the situation model becomes more elaborated, the number of constraints increases, until a single acceptable course of action is implied. Thus, tests for the completeness of an *action* or *plan* can be an important driver of *situation model* construction. The situation model continues to be extended until at least one full course of action, at the level of detail required by current goals, has been generated based on that model.

There is abundant evidence in the cognitive psychology literature that the elaboration of situation models can be driven by the task, i.e., by the goals or plans of the problem solver. Pennington and Hastie (1988) have shown that jurors construct stories to explain evidence in such a way that verdict categories can be mapped onto the stories. Collins, Brown, and Larkin (1980) have shown that proficient readers monitor their comprehension to a degree that is required by the purpose with which they are reading. Voss et al. (1991) showed that experts, but not novices, described an international situation in enough detail to constrain the foreign policy recommendations that they were tasked to provide. Even in verbal object classification, the task influences the level at which objects are categorized (Cruse, 1977). In verbal recall studies, the task determines the depth at which text is encoded (semantic or superficial), and this in turn determines how much is recalled (Craik & Lockhart, 1972).

Klein (1993) has argued that in recognition-primed decision making, courses of action are generated and evaluated one at a time, by contrast with the generation and comparison of multiple options prescribed by analytical models. The most obvious reason for this is the association of situation models with typical responses in procedural processing. But generation and evaluation of single courses of action occurs in knowledge-based processing as well. The reason, according to the R/M model, is that courses of action are not so much retrieved as *designed*, through an iterative process of assessing the situation, extracting constraints on action, evaluating the resulting plan, and reassessing the situation to extract more constraints. Such a process could not be efficiently conducted for more than one plan at a time. Formal models that insist on consideration of multiple options may divert effort from the more productive task of understanding the situation well enough to design a *single* appropriate plan.

Unreliability. A model or plan may be complete but based on shaky data or premises. The second full verification process involves testing for unreliable data or assumptions. Methods used in this step include retracing steps of reasoning while looking for key assumptions that may be weak or unsupported. Mental simulation of a model or plan may reveal that the model can predict actual data, or the plan achieve all goals, only if certain unproved assumptions are made. Verification of reliability may also include a devil's advocate strategy, attempting to generate situations that are contrary to those predicted, alternative interpretations of cues, or alternative means to achieve the same goals. The existence of these alternatives can reveal assumptions underlying the current model or plan.

Virtually any expectation, however certain it may appear, depends on assumptions of one kind or another. As just noted, an effective method for uncovering such assumptions is to image that the expectation is not true, and try to explain how that could be. Figure 9 can serve as an example. Usually, we expect the enemy to place artillery in the region where they plan to attack. Imagining that artillery is placed in a region other than the region of attack can stimulate recognition of hidden assumptions in the usual expectation, i.e., that the artillery is limited in range and in mobility.

There is evidence from the problem-solving literature that experts are more concerned than novices to verify solutions. For example, physics experts use abstract representations of a problem to check their results (Chi, Glaser, & Rees, 1982; Larkin et al., 1980). Experienced physicians were found by Patel and Groen (1991) to spend more time confirming their diagnosis than less experienced doctors. Klein (1993) has proposed that an initial recognitional response to a situation may be subjected to a process of *progressive deepening*, in which it is evaluated and modified if necessary. Progressive deepening can involve tests for completeness or reliability of the course of action. In the context of Naval anti-air warfare decisions (Cohen, 1993), we found that more experienced officers not only waited longer before engaging an unknown contact, but adopted contingency plans (enabling very rapid engagement in case of a hostile act) to mitigate the risk of doing so. Orasanu (1990) found in the commercial air context that proficient air crews were more likely to utilize low workload periods during the cruise phase to prepare contingency plans for anticipated high workload situations.

Conflict. A situation model or plan may be complete and may involve no obvious unreliable data or assumptions. However, a situation model may conflict with observed data, or there may be more than one model that fits the data about as well. Similarly, a plan may fail to satisfy important goals, or there may be more than one plan that satisfies the relevant goals. Another full verification function, therefore, involves discovering conflicts between models and data or between plans and goals, and/or the existence of alternative models or plans. An important method involves generating expectations based on the model or plan, e.g., by mentally simulating future events. New data are compared with these expectations to see if they fit (Noble, 1993). Other methods for uncovering conflict include the devil's advocate strategy described above, explicitly adopting different points of view (e.g., getting into the mind of the enemy, or looking at the situation from the point of view

of a higher echelon or adjacent unit commander), or explicitly asking others for their points of view (e.g., staff members, adjacent, upper, or lower echelon staff or commanders).

Experts may be better than novices in discovering the existence of conflict. In the Naval anti-air warfare context (Cohen, 1993), more experienced officers were better able to generate alternative interpretations of cues regarding target identity or intent.

Facilitation

If no specific problem with the model or plan is identified by either quick verification or full verification, metacognitive processing in the current cycle is complete. But if a specific problem is found, the third major function of metacognition is enlisted: facilitating the construction of an improved model or plan. Whatever the problem that is discovered, three methods are available to solve it:

1. Collecting more data to fill gaps in the model or plan, confirm or disconfirm an assumption, or to resolve conflict
2. Activating existing knowledge in long-term memory, for the same purposes
3. Adding assumptions to fill gaps or resolve conflict, and dropping assumptions when they appear unreliable

Metacognitive processes play a role in choosing among these processes, and in regulating the process that is chosen: (1) in selecting the amount and type of data collection, (2) in directing the search for knowledge in long-term memory, and (3) in adjudicating among competing possible assumptions.

Data collection. Sometimes there is time and opportunity to collect additional data to flesh out or resolve ambiguity in a model or plan, or confirm or disconfirm doubtful assumptions. The decision to collect more data rather than simply think about the problem involves metacognitive judgments regarding the amount of available time, the cost and potential risks of data collection, and the trustworthiness of information sources.

Knowledge activation. Metacognitive processes are crucial in guiding the serial activation of knowledge in long-term memory. This search may be thought of as controlled spreading activation (Lange, 1992). In standard spreading activation, inputs propagate through a network, causing changes in the activation of connected nodes, until the network settles into an equilibrium state. In knowledge-based processing, however, executive processes determine which components of the current model will be attended, thus influencing the portions of long-term memory likely to be activated next (McClelland & Rumelhart, 1986). The values of the attended nodes are fixed, or *clamped*, at a high level of activation (in effect, accepting them provisionally or by assumption) in order to explore their implications. In the next cycle, new nodes may be clamped, and so on, until knowledge is activated that satisfies the goals of the search (or quick verification determines that time has run out). Generic knowledge structures may partially guide this search. For

example, situation assessors may attempt to activate knowledge corresponding to the nodes of a generic enemy plan structure (Figure 2b). Different assessors will frame the situation differently depending on which of these nodes they attend to first. Some may focus attention on knowledge of terrain, others on knowledge of enemy strength, others on knowledge about enemy goals, and others on knowledge of enemy actions.

Metacognitive control may influence search in another way, by determining its *temperature* (Hinton & Sejnowski, 1986), i.e., by adjusting the degree of similarity required for a match between patterns in active memory and stored structures. At high temperatures, the activation net is case wide, and far-fetched ideas have a significant chance of being considered. At low temperatures, an idea must have a very high degree of association with currently active beliefs to have a chance of being activated. High temperatures may be crucial, for example, when all models in the current episodic memory are contradicted by the data, or when no active plan adequately achieves important goals.

Adjusting assumptions. If data collection is infeasible because of limitations in resources, time, or sources of information, and if definitive knowledge is not available or cannot be accessed from long-term memory, the situation assessor may revise his interpretation of the information he has. Metacognitive processes are crucial in the interpretative process of evaluating and revising assumptions.

Assumptions can be defined in two complementary ways (Cohen, 1989):

- a. An assumption is a belief that is not fully or directly supported by evidence.
- b. Assumptions are beliefs that are likely to be retracted in case of conflict with other beliefs.

Absence of direct support (in definition a) can occur for different reasons: The belief may be highly inferential by nature; there may simply be no direct evidence for that type of belief (e.g., inferences about certain elementary particles in physics). Alternatively, direct evidence may be possible in principle but simply not available on this occasion (e.g., assumptions about the reliability of a new sensor or human source of information, or the continued validity of a dated observation). Finally, direct evidence may be possible and available, but simply not yet collected. In connectionist terms, an assumption is a node that becomes activated through indirect links to other activated nodes, rather than direct links to sensory input. (Either direct links to sensory inputs do not exist, or they are not activated on this occasion.)

The second definition of an assumption follows from the first: Because an assumption does not have direct support of its own, it is sensitive to indirect indicators of its validity, such as consonance with other beliefs. In cases of conflict, beliefs with direct support are less likely to be withdrawn than assumptions.

In both senses, being an assumption is a matter of degree (Cohen, 1986). Thus, a conclusion about the enemy's intended location of attack, based on such indirect indicators as the direction of movement of the leading force, would depend relatively heavily on assumptions; a conclusion about intent based on massing of forces and emplacement of artillery would rely less heavily on assumptions; and a conclusion based on the initiation of an artillery barrage and movement of close-in troops in battle formation might be even less assumption-like.

Knowledge-based reasoning often produces a mix of firm beliefs and assumptions. As noted in the section on "Monitoring and regulating cognitive processes," knowledge-based reasoning often involves the activation of information that is only indirectly linked to the current situation. The relevance of this information to the current problem will typically depend on numerous assumptions (more assumptions for high-temperature reasoning, fewer assumptions for low temperature reasoning). Such assumptions might pertain to the similarity of a recalled episode to the current problem, the applicability of a prototype or general explanatory schema, the compatibility with one another of different lines of reasoning that are combined for the first time, and so on. Assumptions of one kind or another are inevitable if decisions are to be made. Knowledge-based reasoning relies in a crucial way on assumptions of this kind in order to fill gaps in situation models and plans.

Decision makers think and act as *if* assumptions were true until there is some reason to doubt them. Conflict between data and a situation model, or between two competing models, provides such a reason for doubt. Conflict indicates that at least one of the beliefs involved in building the models or interpreting the data was false. Conflict may thus trigger a metacognitive process of exposing and questioning assumptions. Other things being equal, the most assumption-like belief, i.e., the least directly and fully supported belief, will be dropped. When more than one assumption is not well supported, other more subtle factors may also play a role. For example, an assumption that is central to a larger variety of important conclusions across a larger range of situations has more indirect support and is more useful; such an assumption is perhaps less likely to be dropped than an assumption that has been adopted on an ad hoc basis for a particular problem. Similarly, an assumption that has conflicted with other beliefs in other situations has more indirect disconfirmation, and may be more likely to be dropped.

The process of revising beliefs to explain conflict requires a variety of metacognitive skills: awareness that conflict exists, an ability to uncover implicit assumptions that have created the conflict, sufficient awareness of the structure of one's beliefs to identify the assumptions that are central to a variety of models and plans, and recall of past episodes in which the same beliefs may have led to a conflict. Finally, the process of assumption revision calls for a balance between the plausibility and the power of the resulting models and plans.

Paths Through the Framework

In this section, we examine how situation assessors operate within the framework of Figure 14. In knowledge-based processing, many different paths

through the situation assessment framework may be followed, depending on the type of problem that is discovered in each cycle and the solution method that is adopted. We will describe a way to analyze any given path in terms of the elementary sequences of cognitive events into which it can be decomposed. Each elementary sequence consists of a single result for verification questions (i.e., an identified problem) and a single choice of facilitation efforts (to solve the problem). Certain sequences are likely to be combined into a path, because of the way that solving one kind of problem can give rise to new problems of a different kind.

Paths through the framework are not explicitly chosen or conscious strategies. Rather they may result from local choices of what to do next. Metacognition involves response to and regulation of other cognitive processes (e.g., memory search, modeling of own or enemy value/action structures, model expansion, analysis, etc.), but the situation assessor need not be able to verbalize either his awareness of the other cognitive processes or the metacognitive processes which monitor and regulate them (Gavelek & Raphael, 1985). In short, metacognition itself can be relatively intuitive and automatic. It may draw on knowledge structures (which contain knowledge about other cognitive processes) which have evolved through long experience in a domain. On the other hand, metacognitive processes can also be relatively analytical, utilizing explicitly taught (or self-taught) methods for verifying and facilitating problem solving.

The three types of problems explored by Full Verification are shown in Figure 16 as three points on a triangle. They represent model or plan incompleteness, unreliable data or assumptions, and the existence of more than one conflicting model or plan. Figure 17a shows how the solution of one of these problems by Facilitation may sometimes lead to the creation of another. The new problem may then be detected and addressed in a subsequent iteration of the Full Verification step. In knowledge-based processing, many different paths through the R/M framework may be followed, depending on the type of problem that is discovered in each cycle and the steps that are taken to correct it.

For example, gaps in an *incomplete* model plan may be filled by making assumptions (shown by the arrow from a to b in Figure 17a), e.g., that the enemy will adopt the worst-case course of action, that a sensor is working as it is supposed to, or that a dated observation is still correct. *Conflict* among different items of evidence may also be resolved by adopting assumptions (shown by the arrow from c to b in Figure 17a), e.g., about the unreliability of one or the other of the conflicting data sources. Assumptions may therefore be justified and necessary in order to arrive at a *complete and coherent* story that explains observed events. Too many such assumptions, however, can lead to trouble. They may blind the decision maker to better hypotheses or plans. In a subsequent verification cycle, the situation assessor (drawing on current episodic memory) may realize that the current model or plan is based on too many *unreliable* assumptions. If he corrects this problem by dropping the unreliable assumptions, the result may again be an *incomplete* model or plan (arrow from b to a), or a set of *conflicting* models and plans (arrow from b to c). The arrows between a and c represent changes in the way conclusions are formulated, without changes in either assumptions or data. Thus, an *incomplete*

model or plan may be fleshed out by listing multiple *conflicting* possibilities; conversely, *conflicting* models or plans may be resolved without assumptions or new data by dropping all but the common elements, i.e., by moving to a more general but *incomplete* model or plan.

In a fundamental sense, all the points inside the triangle of Figure 16 represent the same degree of uncertainty, expressed in different ways (Cohen, Laskey, Vane, McIntyre, and Sak, 1989). The decision maker can move downward in the space toward a single precise belief (i.e., a complete and coherent story or plan) by adopting assumptions, and upward again by dropping assumptions; he can move left or right by selecting a desired level of generality or specificity. But these choices do not change the basic degree of uncertainty. It is the height and width of the triangle, i.e., *the leeway for interpretation*, that represents uncertainty itself in this diagram.

Of course, Facilitation need not always leave uncertainty unchanged. Additional data or knowledge activated from long-term memory, may fill in gaps in an incomplete model or plan, determine the reliability of an assumption, or resolve conflict between competing models and plans. This case is depicted in Figure 17b by shrinking the size of the triangle. The smaller the triangle, the less the total uncertainty, whether it happens to be represented by incompleteness, unreliability, or conflict. New data or firm beliefs diminish a decision maker's freedom to make assumptions, along with the need to do so.

Figures 17a and 17b reflect two sides of problem solving. Adams & Feehrer (1991) summarize the Odyssey curriculum as teaching students how to make problems *simpler*. "Whatever it is, it can be understood" (p. 80), they say, by some combination of interpretation and new information. Figures 17a and 17b represent this duality: the complementary power of assumptions and knowledge.

In the rest of this section, we will describe detailed examples, from actual incidents, of the paths that can arise from combinations of these sequences. Figure 18 provides a key for the symbology in the charts to follow.

Ensuring Plan Completeness and Reliability

As noted above (in the section on "Full verification"), the test for plan completeness is a significant driver of situation model elaboration, until at least one full course of action, at the level of detail required by current goals, has been generated. Klein (1993) has proposed that an initial recognition response to a situation may be subjected to a process of *progressive deepening*, in which it is evaluated and modified if necessary. In the following example, we emphasize: (1) that progressive deepening can involve tests for completeness and reliability, and (2) that the primary vehicle for elaboration or modification of *courses of action* is elaboration of the *situation model* - through activation of additional knowledge.

The path illustrated in Figures 19a and 19b contains the following two elementary sequences, corresponding to cycles 1 and 2, respectively:

1. Verification = incomplete plan; Facilitation = activate LTM knowledge to elaborate situation model, trigger associated action

2. Verification = unreliable assumptions in associated action; Facilitation = activate LTM knowledge to elaborate situation model, trigger actions that form a more reliable plan

Figures 19a and 19b show a series of situation assessment cycles which are typical in progressive deepening. In this example the initial situation model and plan is based on a proactive plan structure. The friendly side has an offensive mission and the goal of planning an attack. They also have a force-size disadvantage. This leads to an intent, on the part of this situation assessor, to increase friendly relative strength by attacking the enemy's center of gravity. Quick verification reveals that this plan is incomplete: Because of this relative unfamiliarity with this enemy, the situation assessor does not know where the enemy's center of gravity is. To facilitate an improved plan, causal models regarding the enemy are activated in order to discover a likely center of gravity. The situation assessor chooses to examine knowledge structures characterizing enemy planning/C² activity. A very simple model of this sort is activated in cycle 2. In this model, the enemy Army commander is represented as making operational plans, and the division commander is represented as carrying them out with little initiative. This immediately suggests that the enemy Army commander is the center of gravity, and this in turn leads to the friendly plan of knocking out the enemy Army commander.

A quick verification of this solution reveals no obvious problems. However, time is available, stakes are high, and this is an unfamiliar situation. Thus, fuller verification is undertaken. The commander and his staff mentally simulate this plan, adopting a devil's advocate approach. Two problems are found. First, the plan may fail in knocking out the Army commander, and second, even if the Army commander is knocked out, the division commander may be able to continue implementing the original plan. To facilitate an improved plan, the causal knowledge structure representing enemy planning/C² activity is expanded. (Such expansion may occur, in connectionist terms, by activating weaker and more distant associations.) The more elaborate enemy planning/C² activity structure includes a new node representing the Army commander's communication both of his plans and of replanning information to the division commander, and another node representing steps taken by the division commander to execute the plan. These new nodes are associated with two new friendly actions: jamming communications between the Army division command post, and preventing the division commander from implementing the plans, e.g., blocking the movement of division troops.

In this example, an initial plan was found to be incomplete and was fleshed out. The fleshed out plan was then found to depend on doubtful assumptions, e.g., regarding the execution of the plan and its outcomes. These doubtful assumptions were then bolstered by adding additional actions to the plan. Each step of improving the plan involved further elaboration of the situation model upon which it was based.

As noted in the section on "Long-term memory", cycle 1 of this example is another illustration of a proactive enemy plan structure. In the earlier example (Figures 5 and 6), the problem was framed in terms of a terrain

pattern corresponding to a kill zone, and an effort was made to influence the enemy's perception of relative strength. In this example, the problem is framed in terms of the enemy's planning/C² activities, and an effort is made to reduce enemy strength. Moreover, in the action of jamming communication of orders from the Army to the division, the assessor proactively influences the division's *goals*.

Testing Expectations and Conflict Resolution

An unreliable situation model may be tested by generating predictions and comparing them to data. However, such tests are not cut-and-dried. Even when data appear to clash with the model, it is possible to find other interpretations of the data that restore consistency and save the model from falsification. The first example in this section represents the following path, corresponding to cycles 1, 2 and 3 in Figures 20a, 20b, and 20c:

1. Verification = unreliable model; Facilitation = collect data to confirm assumptions
2. Verification = conflict between data and model; Facilitation = adjust general assumption about meaning of data to resolve conflict
3. Verification = unfamiliar interpretation of data, find unreliable assumptions underlying data interpretation; Facilitation = adjust specific assumptions underlying data interpretation

In cycle 1 of Figure 20a, a predictive enemy plan structure has been used to infer that the enemy Army will attack in region x. Data in support of this conclusion are somewhat sparse, leading to a metacognitive decision to collect data in order to test this prediction. A temporal plan execution structure of the kind shown in Figure 8 is used to generate the further prediction that the enemy division will move its command post forward in region x. However, the enemy division command post is not in fact observed in region x. In other words, verification by collecting additional data has led to a new problem: conflict.

The conflict is detected by Quick Verification in cycle 2 (Figure 20b). Facilitation now has a choice: It can accept this conflict at face value and drop the belief that the enemy will attack in region x. Alternatively, Facilitation can look for some other explanation of the failure to observe the command post. Facilitation chooses to at least explore the possibility of alternative explanations of the conflicting data.

A new cycle of verification (cycle 3, Figure 20c) focuses on the *validity of the conflicting evidence*: i.e., alternative possible explanations of the failure to observe the command post in region x. In order to generate such explanations, the enemy planning/C² activity structure is elaborated, fleshing out causal connections between the Army intent to attack and the movement of the division command post. This causal structure supports the activation of exceptions, or alternative causal paths, as shown in Figure 20c. Each of these exceptions provides a potential reason for the failure to observe a division command post in region x, even if the intent of the Army is to attack in

region x. Thus, the Army may fail to assign reconnaissance responsibility to the division; instead, the Army may either take over the division or coordinate with the division, in either case performing the reconnaissance itself by air or by train. Even if the division is assigned reconnaissance responsibility, it may decide that the requirement for surprise outweighs the benefits of reconnaissance, and thus drop the reconnaissance mission. Alternatively, the division commander may decide to perform reconnaissance, but attempt to increase the probability of surprise. He may thus move the personnel but not the equipment associated with the command post, or he may decide to take over an already existing forward echelon command post. Finally, the division commander may decide to move both personnel and equipment, but the plan may fail because of destruction of the command post enroute.

The simplest explanation of the failure to observe the command post in region x involves acceptance of the possibility that the command post was destroyed enroute. This exception involves the least disruption of the normal picture of the enemy planning/C² activities. Thus, the Army is regarded as assigning reconnaissance responsibility to the division as usual, the division is regarded as performing the reconnaissance function as usual, and the division is regarded as moving both personnel and equipment in the command post as usual. The break in the normal chain occurs at the last possible step. In this way, the initial "story" represented by the predictive enemy plan structure at the top of Figure 20a is preserved, with only a minor wrinkle in the plot. (Another possibility would be that the enemy successfully moved the command post, but that friendlies failed to observe it because of camouflage, low visibility weather conditions, etc. This was not regarded as a plausible possibility in the present example.) In fact the entire elaborated enemy planning/C² structure shown in Figure 20c was not activated all at once by the situation assessor. Rather, the first exception generated was the one involving the destruction of the command post. Only when the situation assessor was told this was not the case, were further alternatives generated. In general, the pattern of activation of this structure was from the bottom to the top, i.e., starting with minimal disruption of the normal pattern of events and continuing on to increasingly fundamental alterations.

As noted, the example in Figure 20 illustrates how verification of the reliability of a model by collecting new data can lead to the new problem of conflict. In this case, the initial analysis suggested that the enemy would attack in region x while the failure to observe the command post in region x suggested that they would attack elsewhere. The example illustrates further how conflict can be resolved by adopting a new assumption: that the command post was destroyed. The elaboration of the model represented by the new assumption constitutes the simplest and most plausible overall story.

Conflict can also arise at the very earliest stage of the situation assessment process. The next example (cycle 1, Figures 21a, 21b) illustrates this, along with some other processes that may contribute to resolving conflicting data. It contains the following elementary sequences or cycles:

1. Verification = conflicting data; Facilitation = adjust assumptions
2. Verification = conflict between data and model; Facilitation = adjust

general assumption about meaning of data to resolve conflict

3. Verification = unfamiliar interpretation of data, find unreliable assumptions underlying data interpretation; Facilitation = adjust specific assumptions underlying data interpretation
4. Verification = conflict; Facilitation = adjust assumptions
5. Verification = conflict between data and model; Facilitation = adjust general assumption about meaning of data to resolve conflict
6. Verification = unfamiliar interpretation of data, find unreliable assumptions underlying data interpretation; Facilitation = adjust specific assumptions underlying data interpretation
7. Verification = conflict; Facilitation = adjust assumptions
8. Verification = conflict; Facilitation = adjust assumptions
9. Verification = unfamiliar interpretations of data, find unreliable assumptions underlying data interpretations; Facilitation = adjust overall sets of assumptions underlying data interpretations

Cycles 2 and 3, and 5 and 6, illustrate the kind of deliberative conflict resolution already seen in the previous example, in which the decision first decides to assume the evidence doesn't have its usual meaning, and then (in a second cycle) finds a way to make that decision stand. By contrast, cycles 4, 7, and 8 illustrate a more automatic single-cycle process of conflict resolution, in which evidence is simply recognized as having a different meaning. In these cases, the alternative meanings of the evidence are already relatively active in memory and do not need to be searched for. Finally, cycle 9 illustrates an (unfortunately rare) deliberative process of making sure that too much conflicting data have not been explained away.

Figure 21a provides an example in which the initial information gives rise to competing enemy plan structures. The initial evidence can be explained in two ways, just as jurors in Pennington and Hastie's (1988) research may consider competing stories to account for courtroom evidence. In the first story, consideration of enemy principles and goals both point towards an attack in the south: There has been more success in the south, and Soviet doctrine is to exploit success; and the most likely specific objective of the enemy's advance is to take Frankfurt. Considerations of strength are consistent with an attack in the south, i.e., the best supply centers are located in the south. Finally, considerations of location are also consistent with an attack through the south: The terrain in the south provides the best support for armor movement, and the best roads to Frankfurt go through the south.

The alternative story focuses on strength and location. Forces are stronger overall in the north, and the commander in the north is superior to the commander in the south. In terms of location, the forces in the north have better skills at river crossing. These two factors support an attack through the north.

Quick verification reveals that these plan structures are conflicting, and the facilitation process tries to find a coherent explanation that can account for all of the information. Figure 21b shows the result. A single two-tier enemy plan structure has been constructed to coherently account for all the available information. According to the hypothesis generated from this integrated structure, the enemy intends a main attack in the south and a secondary, diversionary attack in the north. The goal of the secondary attack in the north is to reduce friendly strength by drawing it off from the main attack area in the south. The secondary attack thus pursues a subgoal of the main attack, which is to concentrate the enemy's relative strength in the south.

This example is an additional illustration of the predictive use of the enemy plan structure, and a framing of the situation in terms of enemy goals (as in Figures 3 and 4). The situation assessor has concluded that the main attack will be in the south largely because of his focus on enemy goals and principles (see Figure 21a). Recall that we have defined "principles" as a type of fundamental or bedrock value, and "goals" as desired situations for a particular situation. In both strength and location the two candidate stories were approximately equal. But there was no plausible account of enemy goals that supported the conclusion of an attack in the north.

In cycle 2 (Figure 21c) the situation assessor realizes that U.S. reserves are located in the south. This information was available during the initial assessment of the situation, but was not noticed! The situation assessor now concludes that the enemy is not as strong in the south as he thought. Quick verification reveals a conflict. This evidence does not fit the integrated enemy plan structure developed at the conclusion of cycle 1. As in the example of Figure 20, the situation assessor decides to question the conflicting data. In the next cycle (cycle 3, Figure 21c), Quick Verification responds to the need for an unusual interpretation of the strength data, and Full Verification looks for specific unreliable assumptions to back this up. Once again, in order to do so the assessor expands a causal model of the enemy planning/ C^2 activity. The solid arrows in the enemy planning/ C^2 structure (in Figure 21c) represent the normal course of events: The U.S. reserves are located in the south; enemy intelligence observes this fact; and the enemy estimate of its own strength incorporates this fact. The dashed lines represent a possible alternative set of events, i.e., an exception condition. The enemy might not observe U.S. reserves in the south, and thus the enemy might overestimate its relative strength in the south. By accepting the assumption that enemy intelligence has not observed the U.S. reserves in the south, the conflicting information is explained and made consistent with the integrated enemy plan structure. (The situation assessor may have reasoned that if *he* could overlook the location of the U.S. reserves, so could the enemy.)

In this example and in the previous example (Figure 20) we have seen how the attempt to produce a plausible, coherent story accounting for the data can lead to the so-called "confirmation bias." In the confirmation bias, conflicting evidence is reinterpreted to conform with a favored hypothesis (Nisbett & Ross, 1980). Such behavior, however, may be perfectly justifiable in case there are only a small number of outliers or conflicting pieces of

data (Cohen, 1989). The goal of the situation assessor is to produce a single coherent picture of the situation. This can only be achieved if an explanation is found for apparently conflicting data. Moreover, such explanations may well be *true*. Conflict among different lines of reasoning is real evidence that one or more of the assumptions in those lines of reasoning is wrong. Using conflict to identify and correct faulty assumptions can lead to a more accurate knowledge base and to improved situation assessment performance in the future (Cohen, 1986). In fact, we have found that more experienced situation assessors are more, rather than less likely to generate explanations of conflicting data.

But what happens if conflicting data continue to be observed? In this example, that is exactly what happened (Figures 21d, 21e, 21f). In each new case of conflict the situation assessor was able to generate an explanation of the conflicting data that was consistent with the original hypothesis (main attack in the south, diversionary attack in the north). For example, in cycle 4 (Figure 21d) small attacks were observed in the north. This was explained as part of the secondary attack, which was expected in the north. In cycles 5 and 6 (Figure 21e) enemy deep interdiction destroyed bridges in the south, thus hindering any potential advance by the enemy into that sector. The situation assessor proposed two possible explanations: First, the enemy in the south may have had more bridging capability than he had anticipated (his original assessment was that the northern enemy forces had superior bridging capabilities). A second possible explanation is that destruction of the bridges was a mistake. In cycle 7 (Figure 21f) two significant units were observed heading towards the north. The situation assessor again offered two possible explanations: First, this might be a possible feint in support of the diversionary attack, to increase the chance of surprise. Alternatively, this too may have been a mistake. In cycle 8 (Figure 21f) more artillery was observed in the north than in the south. The situation assessor explained this as a possible part of the diversion. He also mentioned the possibility that the artillery possessed longer range than he expected, thus permitting it to strike the south from its location in the north (see Figure 9).

Each of these pieces of conflicting data may be plausibly explained away *if taken by itself*. The problem, of course, is that the process can continue indefinitely. After each piece of conflicting data is explained, the situation assessor may conclude that his favored hypothesis (main attack in the south, diversionary attack in the north) is still supported by *all* the data. Thus each new episode of explaining away appears justified based on the predominance of prior support for the favored hypothesis. However, at some point the accumulation of ad hoc assumptions undermines this justification. It is no longer the case that the predominance of evidence supports the favored hypothesis.

This problem may be detected when the verification process looks for unreliable assumptions, as it does in cycle 9 (Figure 21g). Accurate detection depends on two things happening: The situation assessor must recall the past incidents in which evidence was explained away. And he must ask how many independent explanatory assumptions have been invoked in order to explain away all the conflicting evidence. If too many ad hoc assumptions have been adopted, he may conclude that the favored hypothesis no longer reflects the

most plausible story, and he may decide to explore an alternative. Figure 21g shows the independent explanatory assumptions that were invoked in the context of the integrated predictive enemy plan structure (main attack in south, diversion in north). Three pieces of conflicting evidence - the two units heading north, the observation of more artillery in the north, and the observation of small attacks in the north - can all be explained by the enemy goal of surprise. The failure to take account of the location of U.S. reserves in the south has to be attributed to an error in enemy planning/C². Two possible explanations were proposed for the destruction of bridges in the south. It is either an error in execution or it reflects better capabilities than expected. (Each of these also serves as a possible explanation of other conflicting data). At a minimum, the situation assessor must invoke three separate explanatory principles to account for all the conflicting information, and to retain the integrated plan structure of Figure 21b.

Why did this particular situation assessor continue to explain away conflicting data, never choosing to revisit the possibility of main attack in the north? One possibility is the compelling nature of situation models based on framing the situation in terms of enemy goals. The situation assessor had tried initially, but failed, to find any plausible goal involving enemy attack in the north. Without some way to fill the goal slot in the enemy plan structure supporting an attack in the north, this plan structure remains implausible. (The only plausible goal he did find was to divert opposing force strength in the south.)

A second possible explanation for the failure to change hypotheses involves structural constraints. The realization that too much conflicting evidence has been explained away depends heavily on episodic memory for the current problem. The situation assessor must retain and access a record of the sequential situation models, in which assumptions were adopted about alternative meanings of data. Unfortunately, as noted earlier, retrieval from current episodic memory may be quite weak in a protracted, high workload battlefield scenario.

Reliance on current episodic memory is a form of "reminding" as described by Schank (1982). According to both Schank and our framework, an episode that conflicts with expectations leads to the construction of an explanation. This explanation triggers reminders of previous episodes of failed expectations that were explained *in the same way*. In our framework, this sort of reminding provides reassurance in the process of verifying whether too many unreliable assumptions have been adopted. The new conflict does not cause as much concern because it did not require any (or as much) new elaboration of causal models in order to discover exception conditions. This sort of reminding, which can make explaining away almost automatic, is illustrated in cycles 4, 7, and 8 (Figures 21d and 21f). According to our framework, however, it is also possible to be reminded of previous instances of explaining away that invoked *different* explanatory principles. When this happens there is *less* confidence in the new explanation. It is the latter sort of reminding, unfortunately, that appears more fragile.

Verifying Assumptions and the Reliability of Data

In the previous two sections, we have focused primarily on filling gaps or correcting incompleteness in the situation model or plan, and on resolving conflict. Each of these sections, however, provided examples of verification of the reliability of a model or plan. We saw how filling gaps in a plan could lead to incorporation of unreliable assumptions. We saw how resolving conflict could lead to the adoption of too many ad hoc explanatory assumptions. In both cases the importance of verifying the reliability of those assumptions is clear. In this section, we focus on verification of assumptions and reliability of data per se.

In situation assessment, verification of reliability can occur in a least four ways:

- ! mental simulation to verify the adequacy of a plan (Klein, 1993; Figure 19b);
- ! gathering more data to test a weakly supported model (Noble, 1993; Figure 20a);
- ! expanding causal models to generate alternative interpretations of data (Figures 20c and 21c);
- ! recalling assumptions that have been adopted in the history of a problem-solving session (Figure 21g).

All of these have been illustrated in the examples above. Figure 22 provides another example, in which verification, by keying on a single unreliable assumption, leads to questioning and revision of almost every component of the original situation model and plan. It involves the following elementary sequence:

1. Verification = unfamiliar situation, find conflicts and unreliable assumptions in model and plan; Facilitation = adjust assumptions, activate knowledge

Figure 22a shows an initial predictive enemy plan model with associated actions. The enemy is a guerilla force whose goals are expected to involve embarrassing friendly forces or disrupting communication, canal traffic, and other such activities. Since all of these objectives are in the north, and the enemy is in the south, achieving the objectives requires that the enemy cross a river from south to north. The terrain is mountainous and jungle. The normal procedure in such terrain would be for the enemy to stay off trails in order to avoid ambush. Similarly the enemy would not use the tops of ridges where they would be silhouetted against the sky. These considerations lead to predictions: The enemy will cross the river from south to north and then navigate off the trails and off the ridges toward the northern part of the sector. The normal defensive response would be to defend on high ground, in order to maximize visibility and defensibility.

This reasoning took place prior to formal receipt of mission orders. Thus, time was available to conduct more extensive verification of the initial plan. Mental simulation of the enemy plan and the friendly response involved

elaboration of knowledge structures representing enemy and friendly plan execution. A number of problems were found. A key source of these problems, as shown in the top of Figure 22b, was that the vegetation in that area was higher and thicker than usual. This factor was responsible for activating multiple exceptions, or alternative paths, in the plan execution structures, shown by the dashed lines and boxes. For example, if the enemy used the trails, they would be less susceptible to ambush than expected - since the trails would be hard for the opposing force to locate. By the same token, if the enemy attempted to travel off the trails, they might have considerably more difficulty navigating than expected. Similarly, enemy use of ridges would not lead to silhouetting or skylining, as expected, because of the height and thickness of the vegetation canopy. On the other hand, if they traveled along the contours of the slopes (which would be the normal procedure), they would slide down into the valleys, due to the slippery condition of the slopes. Finally, the friendly response, defending on high ground, makes little sense in this terrain, since visibility would be highly restricted by vegetation. The only defensible location, with high visibility, was the river bank. The result of this verification activity is a revised predictive enemy plan model shown at the bottom of Figure 22b. In this model, it is expected that the enemy will use trails or ridges and that friendlies will defend against the enemy at the river.

Comparing Options and Adjusting Goals

Previous examples in this section have focused primarily on the roles of knowledge structures and metacognitive processes in knowledge-based intuitive processing. A characteristic of this kind of processing is that multiple courses of action are not generated and compared to one another (Klein, 1993). Rather, a single option is activated, verified, modified (if necessary), and possibly rejected. Only then is another course of action activated and verified. Thus, in Figure 19, a single initial plan, to attack the enemy center of gravity, was generated and critiqued; as a result of the critique, it was fleshed out and amplified; but significant alternative plans, or alternative ways of fleshing out and amplifying the plan, were not considered. In Figure 22 as well a single initial plan was generated, to defend on high ground. As a result of verification, it was rejected and replaced by the plan to defend at the river. But the two options were never simultaneously entertained and compared to one another.

Nevertheless, it is sometimes necessary to consider and evaluate multiple options. In some cases, for example, the staff must justify a course of action to the commander. To do so, they argue that it is better than other possibilities, which they must generate for the purpose of the justification. In other cases, there is a genuine disagreement as to the best course of action, e.g., within the staff or between different subordinate units. Under both these kinds of circumstances, knowledge-based analytical behavior can come into play. At the very least, outcomes of the various options may be generated and explicitly compared to one other in terms of goals.

Analytic strategies vary in their formality and systematicity, as suggested by Hammond's (1993) notion of a cognitive continuum between analytical and intuitive behavior. At one extreme, there is a pre-existing general-purpose

method which is explicitly selected and then applied in a pre-determined way to the problem: e.g., generating several qualitatively different options, exhaustively war-gaming each of them, and then constructing a decision matrix that scores all the options on all the evaluative criteria. There is evidence from our own interviews and other sources (e.g., Fallesen, 1993) that analytical strategies of this nature are seldom used. More often, we think, the manner in which options, outcomes, and goals are considered and compared is decided "on the fly," i.e., determined by domain-specific knowledge structures together with local metacognitive choices about the results of earlier steps.

The role of knowledge structures and metacognition is therefore critical in behavior that lies between the extremes of analytical and intuitive processing (Hammond's *quasi-rational* behavior). This behavior (like intuitive behavior) is shaped by specific answers to verification and facilitation steps, rather than arising as an explicit all-or-nothing method.

One example of such a quasi-rational strategy is dominance structuring (Montgomery, 1993). Dominance structuring begins with a tentative choice of a single option, and proceeds to construct a justification of that option as the best (or tied for best). The justification attempts to show that the selected option is as good as or better than all other options with respect to all goals. In the process of constructing this justification, a goal in which the option is not as good as other options may be dropped, the score of the option may be revised on that goal, or that goal may be combined with other goals so that the option turns out to be at least as good as other options on the new, aggregated goal.

Figure 23 provides an example of dominance structuring in the battlefield situation assessment context, and shows how dominance structuring can arise within our framework. It illustrates the key role of intuitive knowledge structures, in addition to general-purpose ones, throughout the process. It also illustrates how the revision of goals is driven not only by the need to justify a favored option, but also by higher level values which those goals are meant to achieve.

The example involves the following elementary sequences, corresponding to cycles in Figure 23:

1. Verification = incomplete plan; Facilitation = collect data
2. Verification = unfamiliar situation, find conflict with goals; Facilitation = adjust assumption in plan to resolve conflict
3. Verification = conflict with other parties; Facilitation = activate knowledge to create evaluation matrix
4. Verification = conflict between matrix and favored option; Facilitation = adjust general assumption that criteria in matrix are valid
5. Verification = unfamiliar interpretation of criteria, find conflict between criteria in matrix and higher-level values and goals; Facilitation = adjust specific assumptions underlying validity of criteria in matrix

6. Verification = unreliable assumption in plan regarding key criterion in new matrix; Facilitation = collect data, adjust assumptions to improve plan

The situation assessor's unit is a heavy, mechanized division with the goal of seizing a town. Figures 23a through 23d show how an initial friendly plan structure for this unit is generated and modified. Two features of the way the situation is initially framed stand out in cycle 1 (Figure 23a): awareness of how the division goal fits into the larger context of corps and theater goals, and the central importance of terrain. A high-level principle guiding his situation assessment behavior was to look first at issues of terrain and mobility, and only secondarily at issues of enemy strength. The reason is that mobility cannot be taken for granted by a heavy unit. It is usually easier to patch up a suitable avenue of approach that encounters too much enemy (e.g., by diverting enemy forces or bolstering own forces) than it is to patch up a plan that avoids the enemy but involves an unsuitable avenue of approach. As a result of this high-level principle, the situation assessor works backward from his division goal (the town to be seized) to high-speed avenues of approach (roads leading to the town) to potential river-crossing sites, and finally to the current division assembly area. He concludes that the river crossings in the north should be used, since they provide immediate access to high-speed roads into the town.

Issues of enemy versus friendly strength enter into planning only in the verification step (cycle 1). The assessor is aware that the plan is incomplete since strength has not been accounted for. In fact, the enemy is more concentrated near the northern river crossings, and this is added to the plan. In cycle 2 (Figure 23a), verification of the more complete plan reveals that casualties from a crossing in that area would be too great. The division might not have sufficient strength left, after such a crossing, to seize the town. Facilitation generates a modified friendly plan (top of Figure 23b): Let another unit secure the crossing sites, and let our division serve as a follow-up force. The follow-up force should encounter little opposition and few casualties, until it reaches the town.

Verification of the modified plan, however, reveals a conflict (cycle 3, Figure 23b). The corps plan specifies crossing the river in the south rather than the north. The other unit prefers to cross the river in the south: The lower concentration of enemy in the south will result both in fewer casualties and in a faster river crossing. In addition, the G-4 also prefers crossing in the south: The lower concentration of enemy will allow the establishment of a logistics base there for follow-on forces.

To understand this conflict better, Facilitation generates a matrix, showing how the two options compare on the relevant evaluative criteria. This is where the process becomes at least in part analytical. The evaluation matrix is a general-purpose knowledge structure, activated by abstract features of this situation: viz., the existence of multiple well-specified alternatives, clearly stated goals or criteria, and a need to justify the preferred course of action.

The evaluation matrix is shown at the bottom of Figure 23b. Evaluation of

the preferred option (crossing the river in the north) in terms of this matrix is highly unsuccessful: The preferred option is worse than the other option on three out of five criteria, and better on only one. In cycle 4 (Figure 23c), however, Verification detects the conflict, and Facilitation responds by deciding to question the assumption that the criteria in the matrix are valid goals. In the next cycle (cycle 5, Figure 23c), Full Verification assesses the criteria in terms of higher-level goals and values. Full Verification asks, What is the basis for these criteria? Do they reflect the real goals in this situation? The answer is no.

The result of verifying each of the original conflicting criteria is shown in Figure 23c. (1) Crossing in the north performs worse in terms of number of casualties expected during the river crossing. The situation assessor now argues that this criterion is unimportant. The other unit can afford casualties, since it has no other missions. The relevant high-level value here is to put the overall mission, and long-term losses, over short-term casualties. (2) The preferred option prevents setting up a logistics base. But the situation assessor also argues that this criterion is unimportant. Setting up a logistics base is not part of the mission statement. (3) The preferred option will result in a slower river crossing because of enemy opposition. But the assessor argues that this criterion should be combined with the other criterion having to do with speed: the time required to get from the river to the town. It is really the overall speed of the operation from the assembly area to the town that matters, not the separate components. Here, the criterion of crossing speed is shown to be inconclusive with respect to the higher level goal of getting to the town quickly.

As a result of this verification process, Facilitation generates a revised evaluation matrix (Figure 23c). The new matrix has three criteria: No interference with real missions (which eliminates other unit casualties and establishment of a logistics base, but retains seizing the town as quickly as possible); high overall speed of operation (which combines speed of river crossing and speed of movement to the town), and sufficient strength to seize the town (one of the original criteria). Evaluation of the preferred option (crossing in the north) in terms of the new evaluation matrix yields a dominance structure: The option is as good as or better than the other option in every respect. Generation of this revised set of criteria leads to corresponding small modifications in the friendly plan structure. For example, the terrain issues now include finding the overall fastest way to the town (including the river crossing) rather than simply the fastest route from the river to the town. Failure to interfere with other missions is added to the issues concerned with interests.

The analytical process has highlighted the importance of speed: Justification of crossing in the north depends on the assumption that its advantage in getting from the river to the town outweighs its disadvantage in crossing the river. Verification of the new plan structure in cycle 6 (Figure 23d) thus focuses on this issue. The assessor realizes he is not as confident as he would like to be in this overall speed advantage (unreliability of assumption). Facilitation takes two steps to strengthen the plan in this regard: Direct access to spot reports regarding enemy locations enables the assessor to make fairly precise estimates of the likely opposition, and thus

the time required to cross the river. Secondly, a deception plan is developed to draw off enemy forces from the northern to the southern crossing sites. With these modifications, the assessor's confidence in the speed advantage of northern crossing is high. The plan was accepted and executed successfully.

On its surface, dominance structuring appears to be a way of rationalizing a choice that has already been made. In this regard, it is highly reminiscent of confirmation bias behavior, in which the interpretation of evidence is revised in order to justify a favored hypothesis (as discussed in the section on "Testing expectations and conflict resolution"). Neither kind of behavior is necessarily wrong, however. Explaining away data may be justified if there is a strong case for a favored hypothesis; it results in a coherent situation picture and, perhaps, a better understanding of what the evidence in fact means. In dominance structuring, the process of revising goals may be justified if a strong enough intuitive case can be made for the initial choice of an option. A decision maker may feel more confidence in his intuitive choice than he does in the inputs to an analytical choice model (i.e., the evaluative criteria that have been articulated). If this is so, he is justified in dropping, revising, or reassessing the criteria in the light of his intuitive choice. This process of modifying goals is a form of *learning*, in which the assessor refines his understanding of his own goals (Cohen, Laskey, & Tolcott, 1987).

The example of Figure 23 shows how reasonable the revision process can be. This assessor did not capriciously or arbitrarily reject a criterion simply because it conflicted with his preferred option. Rather, he used such conflict as a symptom that something might be wrong with the criterion. He then verified the criterion *based on his understanding of the higher-level goals and values relevant to the situation*. The "new" evaluative criteria in fact reflect these goals and values far more closely than the original evaluative criteria did. In that sense, the result of rejecting and combining criteria is *less* arbitrary than the original set of criteria (which emerged rather haphazardly from the discussion with the staff of the other unit and the G-4). Moreover, the new criteria lead to a more persuasive friendly plan, in which interests, strength, and location better reflect the relevant higher-level goals and values (Figure 23d). The persuasiveness of the assessor's case for crossing in the north lies both in its justification of the option he strongly felt to be best, and in its more reasoned relationship to these goals and values.

The analogy between dominance structuring and confirmation bias behavior can be extended one more step. In the confirmation bias case, too many independent instances of explaining away renders the hypothesis suspect. The favored hypothesis may no longer be supported by the preponderance of evidence. In the same way, if justification of an option requires too many *non-motivated* revisions of criteria, combination of criteria into aggregated criteria, or rescoring of options, then the initial choice of a course of action would certainly be cast into doubt. To the extent that revisions can be justified in terms of existing value and goal structures, however, there is little concern. The verification process can serve as a check on unmotivated changes in evaluative criteria. It may ask whether too many criteria or scores were arbitrarily revised in order to justify a particular option. If so, the

facilitation process should try another option. The process may then be iterated with the other option, and again the number of unmotivated revisions required to create a justification may be assessed. The best option, at the end of this process, may be the one that is *most easily* rationalized in terms of intuitive goal and value structures.

Modes of Processing

Throughout our discussion of the interview data, we have drawn on two quite general theoretical distinctions among modes of processing: One distinction is between procedural versus knowledge-based processing, based on the work of Rasmussen (1993). The other distinction is between intuitive and analytic processing, based on the work of Hammond (1993). Although procedural is sometimes equated with intuitive and knowledge-based with analytical, the distinctions are independent. Both data and theory suggest that each is quite useful.

As we have already noted (in the section on "Monitoring and regulating cognitive processes"), procedural processing involves a direct link between the situation, knowledge structures that are activated in that situation, and actions that are associated with those knowledge structures. By contrast, knowledge-based processing requires repeated cycles of processing before action can occur. Typically, it involves the activation of knowledge that is only tenuously connected to the situation, and thus can be reached only through successive stages - such as the expanded causal structures in Figures 20c and 21c, the mental simulation in Figure 22a, or the evaluation matrix in Figure 23b. It may also involve the integration of knowledge that exceeds the capacity of explicit or implicit focus, and which again requires successive stages to integrate (Figure 21g may be an example). Figure 24 informally illustrates the distinction between procedural and knowledge-based processing.

We have already seen how the metacognitive process of Quick Verification helps determine whether processing will be procedural or knowledge-based, as a function of available time, stakes, and confidence in the procedural solution. But metacognition may also play a role in determining whether a solution approach will be analytical or intuitive.

According to Hammond (1993) intuitive processing involves a low degree of conscious awareness and cognitive control, a high rate of data processing, an averaging approach to information integration, and normally distributed errors; it is characterized by high confidence in the answer and low confidence in the method. Analytical processing involves a high degree of control and conscious awareness, slow processing, task-specific modes of information integration, and errors that are small in number but which tend to be large; it is characterized by low confidence in the answer but high confidence in the method.

Hammond claims that task characteristics tend to induce either intuitive or analytical processing. Intuitive processing is induced when inputs are noisy, redundant, simultaneous, numerous, continuous, measured perceptually, and equally important, and when there is no known algorithm or organizing principle for the domain. Analytical processing is induced when inputs are small

in number, nonredundant, discrete, objectively measured, and differentially important and when an algorithm or organizing principle is known. According to Hammond's cognitive continuum hypothesis, intuitive and analytical processing are two ends of a spectrum. Processes may differ in the degree to which they

reflect intuitive or analytical characteristics. Decision makers may also alternate between the two modes: They may feel a need to bolster their confidence in an intuitive solution by an analytical process, or bolster their confidence in an analytical process by an intuitive solution. To some degree, then, the use of an analytical or intuitive method reflects metacognitive judgments by the decision maker regarding confidence in a solution and the most appropriate methods to improve it.

An important distinction for the purposes of training concerns the origin of the long-term memory knowledge structures that are utilized in each case. Intuitive processing tends to involve domain-specific knowledge structures which are developed through experience. Analytical processing tends to involve general-purpose knowledge structures (e.g., evaluation matrices, assessing weights of various factors, algorithmic procedures, etc.) which are developed by explicit instruction, or which have been constructed by reasoning based on rules learned from explicit instruction.

Table 1 provides examples of each of the four major modes of processing (based on combinations of the two distinctions). This taxonomy resembles, but is not identical to, a taxonomy of aircrew decision processes described in Orasanu (1993).

In procedural intuitive processing, (1) preexisting knowledge packages are directly activated by cues in the situation and lead immediately to a response; and (2) the preexisting knowledge packages are based on experience in the domain and are not easily verbalized. This category includes Rasmussen's skill-based behavior, and many instances of Rasmussen's rule-based behavior and Klein's rapid recognition-primed decision making.

In procedural analytical processing, (1) prepackaged knowledge structures are directly activated and are associated with a response; but (2) the origin of the knowledge is instruction rather than domain-specific experience, and the knowledge is often fairly readily verbalized. This category includes other

instances of Rasmussen's rule-based behavior and Klein's rapid recognition-primed decision making. Examples include following doctrinal rules, counting indicators for or against a hypothesis, or simple logic and arithmetic.

Knowledge-based intuitive processing requires repeated activations of domain-specific long-term memory knowledge structures and their integration in working memory. The distinction between procedural and knowledge-based processing (like the distinction between intuitive and analytic) is one of degree. Procedural processing shades off into knowledge-based processing as a function of the number of activation cycles in long-term memory that are required. In general, too, the more cycles of activation required, the more assumptions will be implicit in the final model or plan, and the more potential unreliability. The following examples are on a rough continuum from less to more knowledge-based processing:

- ! Progressive deepening (Klein, 1993), involving the initial *procedural* activation of a situation model and associated response, followed by a set of *knowledge-based* processes. In these processes the response is evaluated by mental simulation of its future consequences; the situation model may be elaborated and fleshed out; and additional constraints on the response may be generated. An example was described in Figure 19.
- ! Explanation-based reasoning (Pennington and Hastie), in which no ready-made situation model is directly activated. Generic schemas specifying what counts as a satisfactory explanation (or story) in the domain are combined with other knowledge to construct explanations of the current situation. We have described many examples of this process. For example, generic enemy plan structures (Figures 2a and 2b) can be combined with enemy goal structures (Figure 4), with terrain structures (Figure 6), or with enemy plan execution structures (Figure 8) in order to construct plan structure explanations of the current situation (Figures 3, 5, and 7, respectively).
- ! Case-based reasoning or reasoning by analogy, in which there is no ready-made generic knowledge structure that fits the situation. Instead, traces of previous episodes are activated that match various unusual aspects of the situation; their similarity to the current situation is assessed; and the associated responses are modified to fit unique aspects of the present situation. Figure 21g illustrates a form of case-based reasoning, in which previous instances of conflict are recalled and examined for similarity or dissimilarity to a current case.
- ! Abduction or exploratory reasoning, which involves the discovery or invention of a hypothesis to explain a novel phenomenon. This is a form of explanation-based reasoning, but there is no ready-to-hand set of knowledge structures (such as those in Figures 3 through 9) from which to construct an explanation. The situation assessor must search long-term memory for appropriate knowledge. An example in science might be the use of the metaphor of fluid flow as the starting point for construction of a theory of electricity.

Metacognitive monitoring and control plays a role in all these cases of knowledge-based intuitive processing: in determining that a procedural

response needs further verification (or recognizing that no adequate procedural response exists); in evaluating intermediate results as new knowledge is activated and combined in long-term memory; in determining the most promising avenues for further exploration; and in identifying the most important assumptions in cases of conflict or unreliability. These functions tend to involve local choices of what to do or think about next, rather than global choices of a solution method.

Knowledge-based analytic processing involves iterated cycles of long-term memory activation in the service of a general-purpose algorithm or solution scheme. Examples include:

- ! decision analysis, in which multiple options must be generated or retrieved and multiple outcomes must be generated or retrieved for each option;
- ! complex logic, in which multiple propositions must be considered and their implications derived, or multiple possible models satisfying the premises must be manipulated (Johnson-Laird, 1983);
- ! other forms of mathematical modeling.

Force ratio calculations and synchronization matrices are common examples of knowledge-based analytic processing in situation assessment.

Metacognition is as important in knowledge-based analytical processing as it is in knowledge-based intuitive processing. Metacognition occurs at almost every stage and involves the same steps of Quick Verification, Full Verification, and Facilitation that support intuitive processing. Metacognitive choices, however, tend to be more global and less local in analytic as compared to intuitive processing:

- ! Decision to analyze the problem and choice of a method. Quick verification may determine that there is plenty of time, that stakes are high, and that there is low confidence in the intuitive solution to the problem. The latter may reflect low confidence in the intuitive method (Hammond, 1993) or the organizational requirement to justify one's conclusions. Facilitation may follow up by associating problem components with elements in a general problem schema (e.g., the objective = select the best course of action; options = x, y, z; goals = A, B, C...). If the components are few in number and can be clearly specified, Facilitation tries to match them to a method that can map the available inputs onto the desired output (e.g., choice of a single option).
- ! Ensuring model completeness. A straightforward verification function is to make sure that all required assessments for each step of the chosen modeling approach are provided (e.g., how well each option scores on each goal). Facilitation helps identify external experts or sources of information that are appropriate for particular inputs (e.g., relevant specialists in fire support, logistics, etc.). Facilitation also identifies parts of one's own long-term memory to explore and activate in order to provide a given input.
- ! Verifying confidence in inputs. Another potential verification function

involves determining that the level of confidence in particular inputs is sufficient relative to their importance to the model. Is a quick and dirty estimate adequate, or is a more refined and accurate one needed? This may involve assessing the sensitivity of results to particular inputs.

- ! Deciding when to stop. The process of analytical modeling repeatedly cycles through the Quick Verification step. If at any time Quick Verification determines that further analysis is not justified (because insufficient time is available, because stakes are not high enough, or because confidence in the solution has risen to an adequate level) then analysis will cease. (Unfortunately, many analytical strategies - unlike intuitive ones - provide no answers at all until they are carried to completion, and then, of course, the answers may be specious.)

- ! Assessing confidence in model conclusions. This Full Verification step involves assessing the results of modeling in terms of: (a) Completeness - Does it satisfy the original task requirements (e.g., justify a single option as best)? (b) Conflict - Does the selected option conform with intuitive results or the results of other modeling approaches? (c) Reliability - Were too many ad hoc assumptions adopted in the modeling approach? If verification fails (and time is available), Facilitation may lead to alternative modeling approaches or else to another iteration of the same modeling approach, in which inputs or parameters are modified.

Some, though not all, proponents of analytical approaches appreciate the key role of these kinds of metacognitive judgments.

Situation Assessment Expertise

The ultimate purpose of the battlefield situation assessment framework is to help identify problems or opportunities in current situation assessment performance, and to construct methods for the improvement of that performance. A first step in that direction involves identifying ways in which more proficient situation assessors differ from less proficient situation assessors. The goal of an improvement technique - whether it is training, decision aids, improved doctrine, or improved personnel selection - will be to promote performance at the level of the most proficient situation assessors.

The discussion addresses three types of expert-novice differences: procedural processing, long-term memory knowledge structures, and metacognitive skills.

Procedural Expert-Novice Differences

Experts will have a larger number of recognitional templates, i.e., relatively direct connections between situation, situation model, and action (Anderson, 1982).

For experts more responses will be automatic, i.e., require little cognitive effort or conscious awareness and control (Schneider & Shiffrin, 1977). Automaticity is distinct from the mere existence of recognitional templates. It requires extensive overpractice in a consistent problem-solving context.

Experts can represent larger amounts of information in working memory by virtue of chunking (Miller, 1967; Newell & Rosenbloom, 1981). Chunking, like automaticity, can result from a large number of consistent experiences in which items of information occur together and thus come to be represented as a single unit. It can also be facilitated by the development of efficient long-term memory representations of relationships or patterns.

Experts will have skilled memory (Ericsson and Polson, 1988), that is, the capacity to associate new information with pre-existing knowledge structures and appropriate retrieval cues, so that it can be immediately activated in relevant situations.

Knowledge-Based Expert-Novice Differences: Long-Term Memory

Experts have more detailed causal models. We have described several knowledge structures that proficient situation assessors use to organize their understanding of the battlefield:

- enemy plans
- enemy goals
- temporal plan execution
- enemy planning/C² activities
- terrain

We have seen how proficient situation assessors can elaborate or expand such causal models in order to fill gaps in their plans or situation models and in order to explain conflicting information.

Experts have better organized knowledge structures. We have seen how proficient situation assessors can frame their understanding of a particular situation in terms of crucial concepts. Such key concepts may include:

- goals (e.g., the need for fuel or POL, the need to exploit success or seize a particular city)
- terrain (e.g., the features of a kill zone, the implications of vegetation growth in a jungle setting, high speed avenues of approach)
- strength (e.g., destroy enemy center of gravity, divert enemy from main attack, funnel enemy into kill zone).

Expert knowledge structures have a larger scope in space and time. Proficient situation assessors utilize knowledge structures that extend beyond those of less proficient decision makers. Thus, proficient decision makers consider the consequences of their own activities for the achievement of higher echelon goals. They focus more attention on the deep battle and on the interests and behaviors of adjacent units.

Experts recall more cases. Proficient situation assessors have a larger repertoire of cases or episodic memories to draw upon in unusual situations. We have seen how such episodic memories can be used to generate exception conditions in the explanation of conflicting data, and may also be used to generate plans in novel situations by reasoning from analogy.

Experts are more likely to frame situations in terms of proactive, rather than predictive or reactive, principles. Several examples of knowledge structures that support proactive performance were examined:

- luring the enemy into a kill zone
- disrupting enemy C²/planning
- diverting the enemy from the main attack

Knowledge-Based Expert-Novice Differences: Metacognitive Skill

There are expert/novice differences in all steps of metacognitive monitoring and regulation: i.e., quick verification, full verification, and facilitation of an improved model.

Quick verification. Proficient situation assessors are more likely to:

- explicitly ask themselves how much time is available before commitment must be made to a decision
- pay explicit attention to the importance of the decision, i.e., the cost of an error, in terms of their own current goals
- explicitly ask themselves how comfortable they are with their understanding of a situation or with the adequacy of a plan.

The result of these skills is that proficient situation assessors will be less likely either to act prematurely or to wait too long to act. Proficient situation assessors will be better at allocating their time and effort among different tasks.

Full verification. Proficient situation assessors will be more likely to:

- search for problems, i.e., critique the current situation model or plan
- use mental simulation to look for gaps in the model or plan
- attempt to generate alternative interpretations of evidence and alternative outcomes of plans in order to expose unreliable assumptions, and to test expectations based on the model or plan

- explicitly note conflicts in the data or conflicts among goals, and to explore other points of view.

As a result, proficient situation assessors will be less likely to produce an overly vague or incomplete situation model; they will be less likely to miss or fail to account for significant data; they will be less likely to overlook unreliable assumptions or conflicts in the data; and they will be less likely to engage in excessive explaining away (confirmation bias).

Facilitation of improved model or plan. Proficient situation assessors will be more likely to:

- select the most appropriate method for correcting a problem in the current situation model or plan, e.g., collect more data, activate additional knowledge in long-term memory, adjust assumptions in the situation model or plan
- have more effective generic knowledge structures to guide search in long-term memory;
- adopt the right threshold for matching knowledge in long-term memory so that the appropriate amount of information is activated
- adopt efficient strategies for searching memory by adopting temporary assumptions in order to explore their implications
- using better judgment in selecting assumptions for adoption and revision.

As a result, proficient situation assessors will have more complete models or plans, a more coherent picture of the situation (including explanations of conflict), and a more plausible total set of assumptions.

Conclusion: Three Approaches to Situation Assessment

Holyoak (1991) has recently distinguished among three generations in the development of cognitive science theories. The first generation is represented by Newell and Simon's General Problem Solver and by decision-analytic normative models of decision making. It focuses on (1) artificial tasks and (2) analytical, general-purpose techniques that involve minimal domain knowledge for their application. Second-generation theories are represented by expert systems and by recognition-primed models of decision making. They focus on (1) real-world tasks with experienced personnel, and (2) highly specialized pattern recognition methods that are heavily dependent on domain knowledge.

The problems with both of these approaches are becoming increasingly well known. First-generation models are too slow; they are incomplete, i.e., they do not address how hypotheses, options, outcomes, and goals are generated; and they are not consistently used by experienced decision makers. Second-generation approaches (which rely on prepackaged knowledge structures) also do not account for how hypotheses, options, outcomes, and goals are generated in relatively novel situations, or how situation assessors handle uncertainty.

Finally, they too are not consistently used by experienced situation assessors. For example, in recent research on Army military planners, Serfaty (1993) found that experienced planners did not perceive more similarities with prior situations compared to novices, did not generate plans more rapidly, tended to see problems as more rather than less complex, were less rather than more confident in their solutions, and felt the need for more rather than less time.

We can perhaps agree with Holyoak that there is an emerging third generation of models. These models account for adaptive, as well as routine, expertise, i.e., the ability to handle novel and uncertain situations. They accommodate significant individual differences in the way problems are solved. And they predict that some skills (but not all) will transfer across tasks. Within our framework situation assessment is a multidimensional skill. It includes both analytical and intuitive methods. And it includes both procedural behavior, with prepackaged knowledge structures, and knowledge-based processing, in which situation models and plans are constructed through an iterative, goal-directed process. We think that models of this kind offer the most promise for the improvement of situation assessment skills.

References

- Adams, M.J., Tenney, Y.J., & Pew, R.W. (1991, December). Strategic workload and the cognitive management of advanced multi-task systems. BBN Systems and Technologies.
- Allport, D.A., Antonis, B., & Reynolds, P. (1972). On the division of attention: A disproof of the single channel hypothesis. Quarterly Journal of Experimental Psychology, 24, 225-235.
- Anderson, J.R. (1982). Acquisition of cognitive skill. Psychological Review, 89(4), 369-406.
- Beach, L.R. (1990). Image theory: Decision making in personal and organizational contexts. NY: John Wiley & Sons.
- Beach, L.R. (1993). Image theory: Personal and organizational decisions. In G.A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsombok (Eds.), Decision making in action: Models and methods (pp. 148-157). Norwood, NJ: Ablex Publishing Corporation.
- Chase, W.G., & Simon, H.A. (1973). Perception in chess. Cognitive Psychology, 4, 55-81.
- Chi, M., Glaser, R., & Rees, E. (1982). Expertise in problem solving. In R.S. Steinberg (Ed.), Advances in the psychology of human intelligence (Vol. 1) (pp. 7-75). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Cohen, M.S. (1986). An expert system framework for non-monotonic reasoning about probabilistic assumptions. In J.F. Lemmer & L.N. Kanal (Eds.), Uncertainty in artificial intelligence. Amsterdam: North-Holland Publishing Co.
- Cohen, M.S. (1989, August). Decision making "biases" and support for assumption-based higher-order reasoning. Proceedings of the Fifth Workshop on Uncertainty and AI. Windsor, Ontario.
- Cohen, M.S. (1993). Metacognitive strategies in support of recognition. To appear in Proceedings, Human Factors and Ergonomics Society. 37th Annual Meeting, Seattle, WA.
- Cohen, M.S., Laskey, K.B., & Tolcott, M.A. (1987, December). A personalized and prescriptive decision aid for choice from a database of options (revised) (Technical Report 87-18). Reston, VA: Decision Science Consortium, Inc.
- Collins, A., Brown, J.S., & Larkin, K.M. (1980). Inference in text understanding. In R.J. Spiro, B.C. Bruce, & W.F. Brewer (Eds.), Theoretical issues in reading comprehension (pp. 385-407). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Connolly, T., & Wagner, W.G. (1988). Decision cycles. In R.L. Cardy, S.M.

- Puffer, & M.M. Newman (Eds.), Advances in information processing in organizations (Vol. 3) (pp. 183-205). Greenwich, CT: JAI Press.
- Craik, F.I.M., & Lockhart, R.S. (1972). Levels of processing: A framework for memory research. Journal of Verbal Learning and Verbal Behavior, 11, 671-684.
- Cruse, D.A. (1977). The pragmatics of lexical specificity. Journal of Linguistics, 13, 153-164.
- Endsley, M.R. (1988). Design and evaluation for situation awareness enhancement. Proceedings of the Human Factors Society 32nd Annual Meeting, 97-101.
- Fallesen, J.J. (1993). Overview of Army tactical planning performance research. Fort Leavenworth, KS: U.S. Army Research Institute.
- Forrest-Pressley, D.L., MacKinnon, G.E., & Waller, T.G. (Eds.) (1985). Metacognition, cognition, and human performance (2 vols.). NY: Academic Press.
- Gavelek, J.R., & Raphael, T.E. (1985). Metacognition, instruction, and the role of questioning activities. In D.L. Forrest-Pressley, G.E. MacKinnon, & T.G. Waller (Eds.), Metacognition, cognition, and human performance (Vol. 2). NY: Academic Press.
- Hammond, K.R. (1993). Naturalistic decision making from a Brunswikian viewpoint: Its past, present, future. In G.A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsombok (Eds.), Decision making in action: Models and methods (pp. 205-227). Norwood, NJ: Ablex Publishing Corporation.
- Hinton, G.E., & Sejnowski, T.J. (1986). Learning and relearning in Boltzmann machines. In J.L. McClelland, D.E. Rumelhart, & the PDP Research Group (Eds.), Parallel distributed processing: Explorations in the microstructure of cognition: Volume 1: Foundations (pp. 282-317). Cambridge, MA: The MIT Press.
- Holyoak, K.J. (1991). Symbolic connectionism: Toward third-generation theories of expertise. In K.A. Ericsson & J. Smith (Eds.), Toward a general theory of expertise. Cambridge: Cambridge University Press.
- Johnson-Laird, P.N. (1983). Mental models. Cambridge, MA: Harvard University Press.
- Kahan, J.P., Worley, D.R., & Stasz, C. (1989). Understanding commanders' information needs. Santa Monica, CA: The RAND Corporation.
- Klein, G.A. (1993). A recognition-primed decision (RPD) model of rapid decision making. In G.A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsombok (Eds.), Decision making in action: Models and methods (pp. 138-147). Norwood, NJ: Ablex Publishing Corporation.
- Kosslyn, S.M., & Koenig, O. (1992). Wet mind: The new cognitive neuroscience

NY: The Free Press.

- Lange, T.E. (1992). Hybrid connectionist models: Temporary bridges over the gap between the symbolic and the subsymbolic. In J. Dinsmore (Ed.), The symbolic and connectionist paradigms closing the gap (pp. 237-290). Hilldale, NJ: Lawrence Erlbaum Associates.
- Larkin, J., McDermott, J., Simon, D.P., & Simon, H.A. (1980). Expert and novice performance in solving physics problems. Science, 20(208), 1335-1342.
- McClelland, J.L., & Rumelhart, D.E. (1986). A distributed model of human learning and memory. In J.L. McClelland, D.E. Rumelhart, & the PDP Research Group (Eds.), Parallel distributed processing: Explorations in the micro-structure of cognition: Volume 2: Psychological and biological models (pp. 170-215). Cambridge, MA: The MIT Press.
- Miller, G.A. (1967). The magical number seven, plus or minus two: Some limits on our capacity for processing information (Chapter 2). In G.A. Miller (Ed.), The psychology of communication: Seven essays (pp. 14-44). Baltimore, MD: Penguin Books, Inc.
- Montgomery, H. (1993). The search for a dominance structure in decision making: Examining the evidence. In G.A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsombok (Eds.), Decision making in action: Models and methods (pp. 182-187). Norwood, NJ: Ablex Publishing Corporation.
- Neisser, U. (1976). Cognition and reality. Principles and implications of cognitive psychology San Francisco: W.H. Freeman and Company.
- Newell, A., & Rosenbloom, P.S. (1981). Mechanisms of skill acquisition and the law of practice. In J.R. Anderson (Ed.), Cognitive skills and their acquisition (pp. 1-55). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Nisbett, R., & Ross, L. (1980). Human inference: Strategies and shortcomings of social judgment. Englewood Cliffs, NJ: Prentice-Hall, Inc.
- Noble, D. (1993). A model to support development of situation assessment aids. In G.A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsombok (Eds.), Decision making in action: Models and methods. Norwood, NJ: Ablex Publishing Corporation.
- Patel, V.L., & Groen, G.J. (1991). The general and specific nature of medical expertise: A critical look. In K.A. Ericsson & J. Smith (Eds.), Toward a general theory of expertise: Prospects and limits. Cambridge: Cambridge University Press.
- Pearl, J. (1988). Probabilistic reasoning in intelligence systems: Networks of plausible inference. San Mateo, CA: Morgan Kaufmann Publishers, Inc.
- Pennington, N., & Hastie, R. (1988). Explanation-based decision making: Effects of memory structure on judgment. Journal of Experimental Psychol-

ogy: Learning, Memory, & Cognition, 14(3), 521-533.

- Rasmussen, J. (1993). Deciding and doing: Decision making in natural contexts. In G.A. Klein, J. Orasanu, R. Calderwood, & C.E. Zsombok (Eds.), Decision making in action: Models and methods (pp 158-171). Norwood, NJ: Ablex Publishing Corporation.
- Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986). Learning internal representations by error propagation. In D.E. Rumelhart, J.L. McClelland, & the PDP Research Group (Eds.), Parallel distributed processing: Explorations in the microstructure of cognition: Volume 1: Foundations (pp. 318-362). Cambridge, MA: The MIT Press.
- Rumelhart, D.E., Smolensky, P., McClelland, J.L., & Hinton, G.E. (1986). Schemata and sequential thought processes in PDP models. In J.L. McClelland, D.E. Rumelhart, & the PDP Research Group (Eds.), Parallel distributed processing: Explorations in the microstructure of cognition: Volume 2: Psychological and biological models (pp. 7-57). Cambridge, MA: The MIT Press.
- Sanford, A.J., & Garrod, S.C. (1981). Understanding written language. NY: Wiley.
- Schank, R.C. (1982). Dynamic memory. Cambridge: Cambridge University Press.
- Schneider, W., & Shiffrin, R.M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. Psychological Review, 84, 1-66.
- Serfaty, D. (1993, February 4-5). Hypotheses and recent findings on command decision-making expertise. Presentation at U.S. Army Research Institute Workshop on Developing Expertise in Command Decision-Making, Fort Leavenworth, Kansas.
- Shiffrin, R.M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. Psychological Review, 84, 127-190.

Voss, J.F., Wolfe, C.R., Lawrence, J.A., & Engle, R.A. (1991). From representation to decision: An analysis of problem solving in international relations. In R.J. Sternberg & P.A. Frensch (Eds.), Complex problem solving: Principles and mechanisms. Hillsdale, NJ: Lawrence Erlbaum Associates.